arXiv:1703.09945v1 [math.GR] 29 Mar 2017

ON THE CONNECTIVITY OF LEVEL SETS OF AUTOMORPHISMS OF FREE GROUPS, WITH APPLICATIONS TO DECISION PROBLEMS

STEFANO FRANCAVIGLIA

Dipartimento di Matematica of the University of Bologna

ARMANDO MARTINO

Mathematical Sciences, University of Southampton

Abstract. We show that the level sets of automorphisms of free groups with respect to the Lipschitz metric are connected as subsets of Culler-Vogtmann space. In fact we prove our result in a more general setting of deformation spaces. As applications, we give metric solutions of the conjugacy problem for irreducible automorphisms and the detection of reducibility. We additionally prove technical results that may be of independent interest — such as the fact that the set of displacements is well ordered.

Contents

1. Introduction 2
2. Preliminaries 6
2.1. Motivation for new definitions 6
2.2. General definitions and notations 7
2.3. Simplicial structure 9
2.4. Simplicial bordification 10
2.5. Horospheres and regeneration 11
2.6. The groups Aut(Γ) and Out(Γ) 12
3. PL-maps, gate structures, and optimal maps. 13
3.1. PL-maps 13
4. Displacement function and train track maps for automorphisms 23
5. Behaviour of λ at bordification points 30
6. Convexity properties of the displacement function 36

E-mail addresses: stefano.francaviglia@unibo.it, A.Martino@soton.ac.uk.

1991 *Mathematics Subject Classification.* 20E06, 20E36, 20E08.

7. Existence of minimal displaced points and train tracks at the bordification 39
8. Statement of main theorem and regeneration of paths in the bordification 47
9. Preparation to peak reduction 50
10. The end of the proof of Theorem 8.3: peak reduction on simplicial paths 54
11. Applications 57
12. Appendix: proof of Theorem 9.5 60
References 65

1. Introduction

We consider F_n the free group of rank n, usually with a basis B (a free generating set). We are interested in the automorphism group, $\mathrm{Aut}(F_n)$ and the Outer automorphism group, which is defined as $\mathrm{Out}(F_n) = \mathrm{Aut}(F_n)/\mathrm{Inn}(F_n)$.

In recent years there has been a great deal of attention given to the Lipschitz metric on CV_n, Culler-Vogtmann space, see [1], [2], [3] for instance. It has been considered even more generally in [16].

The main goal of the paper, Theorem 8.3, is to prove a result about the connectedness of the level sets of the displacement function λ_ϕ of $\phi \in \mathrm{Out}(F_n)$. That is, one considers the Lipschitz metric then one can take the infimum of all displacements of points for ϕ in CV_n. This infimum may not be realised in general. However, we show that for any $\epsilon > 0$, the set of all points of CV_n displaced at most ϵ more than this infimum is connected. Formally[1],

Theorem 8.3 (Level sets are connected)**.** Let $\phi \in \mathrm{Out}(\Gamma)$. For any $\varepsilon > 0$ the set

$$\{X \in \mathcal{O}(\Gamma) : \lambda_\phi(X) \leq \lambda(\phi) + \varepsilon\}$$

is connected in $\mathcal{O}(\Gamma)$ by simplicial paths. The set

$$\{X \in \overline{\mathcal{O}(\Gamma)}^{\infty} : \lambda_\phi(X) = \lambda(\phi)\}$$

is connected by simplicial paths in $\overline{\mathcal{O}(\Gamma)}^{\infty}$.

This result is of independent interest, and surprisingly strong. We also show how to deduce algorithmic results from this geometric one. We note that these algorithmic results are already known, but are mainly a demonstration of the power of the result. In general, if one wants to check a property P, detectable by simplicial maps somewhere in outer space, an algorithm of the type "go to neighbour simplex

[1]Theorem 8.3 holds in a general setting, the result in CV_n is for $\Gamma = F_n$. Here $\overline{\mathcal{O}(\Gamma)}^{\infty}$ denotes the simplicial bordification of the outer space. Precise definitions are given through the paper.

and check P" clearly stops if it finds a points having P, but it has no a stopping procedure. Connectivity of level sets provides stopping criteria.

Specifically, we solve the conjugacy problem for irreducible automorphisms and prove that it is determining whether an automorphism is irreducible or not is decidable. The conjugacy problem for irreducible automorphisms has already been solved by [14] and [15]. Deciding irreducibility of automorphisms has been proven by [13] and improved in [12]. While our solution of the former has some similarities to that in [14] (namely, peak reduction), our approach is distinctly geometric and provides a uniform framework for dealing with this type of problems for general deformation spaces. Moreover, our connectivity result is general and does not assume irreducibility, so there is hope of pushing these techniques further even though one generally needs irreducibility in order to avoid singularities (which for us means entering the thin part of CV_n).

In proving Theorem 8.3 we obtain a collection of results that may be of independent interest. Namely:

- In Section 6 we give a detailed analysis of the convexity properties of the displacement function.
- We prove that the global simplex-displacement spectrum of $\mathrm{Aut}(F_n)$ is well-ordered. (Theorem 7.2.)
- Generalising a result of [5], we show that local minima of the displacement are global minima (Lemma 4.19). This allows, together with Theorem 7.2 to implement an efficient gradient method for finding train tracks.
- We study the behaviour of the displacement at bordification points, providing a characterisation of those points where the displacement does not jump. (Corollary 7.8.)
- We show that train tracks at infinity minimise the displacement. (Therem 7.11.)
- Given an automorphism ϕ, we show that any invariant free factor is visible in a train track map. (Corollaries 7.12 and 7.13)
- We also wish to mention Theorem 9.5: a technical result, with explicit estimates, that can be phrased as "Folds of illegal turns in a simplex may be closely read in close simplices".

The main results of the paper are proved by induction on the rank, and Theorem 8.3 is assumed inductively true in many points.

The paper is extremely technical, even though the ideas are fundamentally straightforward. In order to motivate the detailed discussion, we provide here the two algorithms for solving conjugacy in the irreducible case and for detecting irreducibility. We present these algorithms as naively as possible, in order to make them more accessible.

That is, one could understand and implement them without any knowledge of the Lipschitz metric, Culler-Vogtmann space or train track maps. As such we have made no attempt to streamline the algorithms in any way; they are brute force searches in an exponential space.

However, we would stress that our point of view is fundamentally that these procedures would be better run as path searches in Culler-Vogtmann space, enumerating optimal PL-maps and calculating displacements via candidates. That abundance of terminology would make the algorithms much harder to describe, so we instead translate everything to a more manageable setting; bases of F_n and generating sets for $\mathrm{Out}(F_n)$. However, the technical point of view is more helpful in developing an intuition of the processes and is likely the way to vastly improve the algorithmic complexity.

Let us know describe our algorithms, whose correctness is proved at the end of the paper. First, we recall some terminology. In order to work algorithmically with $\mathrm{Out}(F_n)$ we need a generating set. The best known of these is the set of Nielsen generators, but it is more convenient for us to work with the following:

Definition 1.1 (CMT Automorphisms, [9] and [8])**.** A *CMT* automorphism of F_n is one that is induced by a change of maximal tree. More precisely, let X be a graph with fundamental group of rank n, and let R be the rose of rank n (the graph with one vertex and n edges). Let T, T' be two maximal trees of X, and let $\rho_T, \rho_{T'}$ be the corresponding projections from X to R. Then the (outer) automorphism induced by changing the maximal tree from T to T' is the (homotopy class of the) map $\rho_{T'}{\rho_T}^{-1}$, where the inverse denotes a homotopy inverse.

The set of CMT maps includes all Whitehead automorphisms, (see [9], Theorem 5.5 and [18]) and is a finite set which generates $\mathrm{Out}(F_n)$.

For convenience, we will include all graph automorphisms of R, including inversions of generators, in the set of CMT automorphisms.

Next we need a notion of size of an automorphism, which will provide a termination criterion for our algorithms.

Definition 1.2. Let $\phi \in \mathrm{Out}(F_n)$, and let B be a basis of F_n. Define $||\phi||_B$ to be $\sup_{1\neq g\in F_n} \frac{||\phi g||_B}{||g||_B}$, where $||g||_B$ denotes the cyclic reduced length of g with respect to B. This supremum is a maximum and is realised by an element of cyclic length ≤ 2.

Remark 1.3. Note that for any constant, C, there are only finitely many $\phi \in \mathrm{Out}(F_n)$ such that $||\phi||_B \leq C$.

Our first application is then as follows. (See Section 11 for the proof.)

Theorem 1.4. *The following is an algorithm to determine whether two irreducible automorphisms are conjugate.*

Let ϕ, ψ be two irreducible outer automorphisms of F_n, and B a basis of F_n.

- *Choose any $\mu > \max\{||\phi||_B, ||\psi||_B\}$.*
- *Inductively construct a finite set, $S = S_{\phi,\mu}$, as follows (which depends on both ϕ and μ):*
  - *Start with $S_0 = \{\phi\}$.*
  - *Set $K = n(3n-3)\mu^{3n-1}$.*
  - *Inductively put S_{i+1} to be all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S_i, ζ is any CMT automorphism, subject to the constraint that $||\zeta\phi_i\zeta^{-1}||_B \leq K$. (We include the identity as a CMT automorphism so that $S_{i-1} \subseteq S_i$).*
  - *End this process when $S_i = S_{i+1}$, and let this final set be S.*
- *Then ψ is conjugate to ϕ if and only if $\psi \in S$.*

Of course, one would like to also be able to decide when an automorphism is irreducible when it is given by images of a basis, for instance. In order to do so, we recall the definition of irreducibility.

Definition 1.5 (see [4])**.** An (outer) automorphism, ψ of F_n is called *reducible* if there are free factors, $F_1, \ldots, F_k, F_\infty$ such that $F_n = F_1 * \ldots F_k * F_\infty$ and each $\psi(F_i)$ is conjugate to F_{i+1} (subscripts taken modulo k). If $k = 1$ we further require that $F_\infty \neq 1$. (In general $\psi(F_\infty)$ is not conjugate to F_∞. Otherwise ϕ is called irreducible.

Equivalently, ψ is reducible if it is represented by a homotopy equivalence, f, on a core graph, X, such that X has a proper, homotopically non-trivial subgraph, X_0, such that $f(X_0) = X_0$. (Being represented by f means that there is an isomorphism, $\tau : F_n \to \pi_1(X)$ such that $\psi = \tau^{-1} f_* \tau$).

We add the following, which constitutes an obvious way that one can detect irreducibility by inspection.

Definition 1.6. Consider F_n with basis B and let ψ be an outer automorphism of F_n. We say that ψ is visibly reducible with respect to B, or simply visibly reducible, if there exist disjoint subsets $B_1, \ldots, B_k$ of B such that $\psi(\langle B_i \rangle)$ is conjugate to $\langle B_{i+1} \rangle$ (with subscripts taken modulo k). If $k = 1$ we also require that $B_1 \neq B$.

More generally, we say that a homotopy equivalence on the rose is visibly reducible if it is visibly reducible with respect to the basis given by the edges of the rose.

This is, in fact, easy to check by classical methods due to Stallings, [17].

Lemma 1.7. *If ψ is visibly reducible, it is reducible. Moreover, there is an algorithm to determine if ψ is visibly reducible with respect to B.*

Proof. The first statement is clear, since each subset of a basis generates a free factor, and disjoint subsets generate complementary free factors. Since there are only finitely many subsets to check, we simply need to determine if the conditions that $\psi(\langle B_i \rangle)$ is conjugate to $\langle B_{i+1} \rangle$ hold. But this can readily be checked since two subgroups of a free group are conjugate if and only if the core of their Stallings graphs are equal, [17]. □

We can now describe our second algorithm. (See Section 11 for the proof.)

Theorem 1.8. *The following is an algorithm to determine whether or not an outer automorphism of F_n is irreducible.*

Let ϕ be an automorphism of F_n, and B a basis of F_n. Construct $S = S_\phi$ as above. Namely,

- *Choose any $\mu > ||\phi||_B$.*
- *Inductively construct the finite set, $S = S_{\phi,\mu}$:*
  - *Start with $S_0 = \{\phi\}$.*
  - *Set $K = n(3n-3)\mu^{3n-1}$.*
  - *Inductively put S_{i+1} to be all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S_i, ζ is any CMT automorphism, subject to the constraint that $||\zeta\phi_i\zeta^{-1}||_B \leq K$. (We include the identity as a CMT automorphism so that $S_{i-1} \subseteq S_i$).*
  - *End this process when $S_i = S_{i+1}$, and let this final set be S.*
- *Let S^+ be the set of all possible automorphisms $\zeta\phi_i\zeta^{-1}$, where ϕ_i is any element of S, ζ is any CMT automorphism, with no other constraint.*
- *If some $\psi \in S^+$ is visibly reducible with respect to B, then ϕ is reducible. Otherwise, ϕ is irreducible.*

2. Preliminaries

2.1. **Motivation for new definitions.** First, we want to motivate the definitions that we are going to introduce. This is because they are a little different and at times more complicated than those usually present in literature. Our aim is to study automorphisms of free groups which are possibly reducible. If Γ is a marked graph with $\pi_1(\Gamma) = F$ a free group, and $\phi \in \mathrm{Aut}(F)$, then ϕ can be represented by a simplicial map (sending vertices to vertices and edges to edge paths) $f : \Gamma \to \Gamma$. That is, f represents ϕ if there is an isomorphism $\tau : F_n \to \pi_1(\Gamma)$ such that $\phi = \tau^{-1} f_* \tau$.

If ϕ is reducible it may happen that there is a collection of disjoint connected sub-graphs $\Gamma_1, \ldots, \Gamma_k$ such that f permutes the Γ_i's. In order to study the properties of ϕ it may help to collapse such an invariant collection. If we want to keep track of all the relevant information,

we will be faced with the study of some particular kind of moduli space. Namely, moduli spaces of actions on trees with possibly non-trivial vertex stabilizers (when we collapse the Γ_i's) and product of such spaces (when we consider the restriction to ϕ to the Γ_i's.)

Since the notation that we are going to introduce may be cumbersome, we will often abuse it, making no distinction between an element of outer spaces and its (projective) class, or between G-trees and G-graphs.

2.2. **General definitions and notations.** Let $G = G_1 * \cdots * G_p * F_n$ be a free product of groups, where F_n denotes the free group of rank n (we allow n to be zero, in that case we omit F_n). We do not assume the G_i's are indecomposable. Throughout the paper, G will be a free group. In particular, each G_i will be a free factor of G. Thus, there is no uniqueness of this free product decomposition, since G has many different splittings as a free product. We use the notation $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$ to indicate a splitting of G. We briefly recall the definition of the outer space $\mathcal{O}(G)$ of G corresponding to the splitting $\mathcal{G}$, referring to [5, 11] for a detailed discussion of definitions and general properties of $\mathcal{O}(G)$.

Definition 2.1 (Outer space)**.** The (projectivized) outer space of G, relative to the splitting $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$, consists of (projective) classes of minimal simplicial metric G-trees X such that:

- For every G_i there is exactly one orbit of vertices whose stabilizer is conjugate to G_i. Such vertices are called *non-free*. Remaining vertices have trivial stabilizer and are called *free* vertices.
- X has no redundant vertex (i.e. free and two-valent).
- X has trivial edge stabilizers.

We use the notation $\mathcal{O}(G)$ to indicate the outer space of G and, if we want to emphasize the splitting, we write $\mathcal{O}(G; \mathcal{G})$ (and $\mathbb{P}\mathcal{O}(G)$ and $\mathbb{P}\mathcal{O}(G; \mathcal{G})$ for projectivized ones). We stress here that when the distinction between $\mathcal{O}(G)$ and $\mathbb{P}\mathcal{O}(G)$ is not crucial, we will often make no distinction between $\mathcal{O}(G)$ and $\mathbb{P}\mathcal{O}(G)$.

Remark 2.2. The equivalence relation that defines $\mathbb{P}\mathcal{O}(G)$ is the following: X and Y are equivalent if there is an homothety (isometry plus a rescaling by a positive number) $X \to Y$ conjugating the actions of G on X and Y. If there is no ambiguity, we will make no distinction between a G-tree X, its class in $\mathcal{O}(G)$, and its projective class in $\mathbb{P}\mathcal{O}(G)$.

Remark 2.3. If $G = G_1$, then $\mathcal{O}(G)$ consists of a single element: a point stabilized by G_1, and in this case the equivalence relation is trivial.

Definition 2.4. A splitting $\mathcal{S} : G = H_1 * \ldots ... * H_q * F_r$ is a *sub-splitting* of $G = G_1 * \cdots * G_p * F_n$ if each H_i decomposes as

$$H_i = G_{i_1} * \ldots G_{i_l} * F_s$$

where F_s is a free factor of F_n and $i_1, \ldots, i_l \in \{1, \ldots, p\}$. Sometime we will make use of the notation $\mathcal{O}(G; H_1 * \cdots * H_q * F_r)$ to mean $\mathcal{O}(G; \mathcal{S})$.

Remark 2.5. If $T \in \mathcal{O}(G)$, the quotient $X = G\backslash T$ is a finite metric graph of groups with trivial edge-groups, together with a marking that identifies G with $\pi_i(X)$ and maps the G_i's to the vertex-groups. We will refer to such graphs as G-graphs (or $(G, \mathcal{G})$-graph if we need to specify the splitting). On the other hand, given a metric G-graph Y, its universal cover $\widetilde{Y}$ is a G-tree in the unprojectivized outer space $\mathcal{O}(G)$. Here $\pi_1(X)$ means the fundamental group of X as graph of groups. (The fundamental group of X as a topological space is just F_n.)

Notation 2.6. If there is no ambiguity we will make no distinction between G-tress and G-graphs. In case of necessity we will use the *tilde*-notation: X for a G-graph and $\widetilde{X}$ for a G-tree, meaning that $X = G\backslash\widetilde{X}$. As usual, if $x \in X$ then $\widetilde{x}$ will denote a lift of x in $\widetilde{X}$. The same for subsets: if $A \subset X$ then $\widetilde{A} \subset \widetilde{X}$ is one of its lifts.

Definition 2.7 (Immersed loops)**.** A path γ in a G-graph X is called *immersed* if it is has a lift $\widetilde{\gamma}$ in $\widetilde{X}$ which is embedded. (Note that γ could not be topologically immersed.)

Let X be a G-graph and let $\Gamma = \sqcup_i \Gamma_i$ be a sub-graph of X whose connected components Γ_i have non-trivial fundamental groups (as graphs of groups). Then Γ induces a sub-splitting $\mathcal{S}$ of $\mathcal{G}$ where the factor-groups H_j are either

- the fundamental groups $\pi_1(\Gamma_i)$, or
- the vertex-groups of non-free vertices in $X \setminus \Gamma$.

Notation 2.8. We will use the notation

$$\mathcal{O}(X) := \mathcal{O}(G) \qquad \mathcal{O}(X/\Gamma) := \mathcal{O}(G; \mathcal{S}) \qquad \mathcal{O}(\Gamma) := \Pi_i \mathcal{O}(\pi_1(\Gamma_i))$$

The above notation leads to the following general definition.

Definition 2.9. Let $\Gamma = \sqcup_{i=1}^k \Gamma_i$ be a finite disjoint union of finite connected graphs of groups Γ_i with trivial edge-groups and non-trivial fundamental group (as graphs of groups). Let $H_i = \pi_1(\Gamma_i)$, equipped with the splitting given by the vertex-groups of Γ_i (hence Γ_i is an H_i-graph). We define $\mathcal{O}(\Gamma)$ as the product of the $\mathcal{O}(H_i)$'s

$$\mathcal{O}(\Gamma) = \Pi_{i=1}^k \mathcal{O}(H_i) = \Pi_{i=1}^k \mathcal{O}(\Gamma_i).$$

We tacitly identify $X = (X_1, \ldots, X_k) \in \mathcal{O}(\Gamma)$ with the labelled disjoint union $X = \sqcup_i X_i$. An element of $\mathcal{O}(\Gamma)$ will be also called **Γ-graph** (or **Γ-tree** if we work with universal covers).

Here we need to be more precise about projectivization. There is a natural action of $\mathbb{R}^+$ on $\mathcal{O}(\Gamma)$ given by scaling each component by the same amount. The quotient of $\mathcal{O}(\Gamma)$ by such action is the projective outer space of Γ and it is denoted by $\mathbb{P}\mathcal{O}(\Gamma)$.

Notation 2.10. In what follows we use the following convention:

- G will always mean a group with a splitting $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$;
- $\Gamma = \sqcup\Gamma_i$ will always mean that Γ is a finite disjoint union of finite graphs of groups Γ_i, each with trivial edge-groups and non-trivial fundamental group $H_i = \pi_i(\Gamma_i)$, each H_i being equipped with the splitting given by the vertex-groups.

Definition 2.11. The rank of the splitting $G = G_1 * \cdots * G_p * F_n$ is $n + p$. The rank of a graph of groups X is the rank of the splitting induced on $\pi_1(X)$, finally if $\Gamma = \sqcup\Gamma_i$ we set

$$\operatorname{rank}(\Gamma) = \sum_i \operatorname{rank}(\Gamma_i).$$

By definition, the rank is a natural number greater or equal to one. Note the the rank of a graph of groups X is nothing but the rank of its fundamental group as a topological space plus the number of non-free vertices.

We will also consider moduli spaces with marked points.

Notation 2.12. The moduli space of G-trees with k labelled points $p_1, \dots, p_k$ (not necessarily distinct) is denoted by $\mathcal{O}(G, k)$ or $\mathcal{O}(G; \mathcal{G}, k)$. If $\Gamma = \sqcup_{i=1}^s \Gamma_i$, given $k_1, \dots, k_s$ we set

$$\mathcal{O}(\Gamma, k_1, \dots, k_s) = \Pi_i \mathcal{O}(\Gamma_i, k_i).$$

If X is a Γ-graph and $A \subset X$ is a subgraph whose components have non-trivial fundamental group, we define $\mathcal{O}(X/A)$ and $\mathcal{O}(A)$ as in Notation 2.8.

2.3. **Simplicial structure.**

Definition 2.13 (Open simplices)**.** Given a G-tree X, the open simplex Δ_X is the set of G-trees equivariantly homeomorphic to X . The Euclidean topology on Δ_X is given by assigning a G-invariant positive lent $L_X(e)$ to each edge e of X. Therefore, if X has k orbit of edges, then Δ_X is isomorphic to the standard open $(k-1)$-simplex if we work in $\mathbb{P}\mathcal{O}(G)$, and to the positive cone over it if we work on $\mathcal{O}(G)$. Given two elements X, Y in the same simplex $\Delta \subset \mathcal{O}(G)$ we define the **Euclidean** sup-distance $d_\Delta^{Euclid}(X, Y)$ ($d_\Delta(X, Y)$ for short)

$$d_\Delta^{Euclid}(X, Y) = d_\Delta(X, Y) = \max_{e \text{ edge}} |L_X(e) - L_Y(e)|.$$

Such definitions extend to the case of $\Gamma = \sqcup_i \Gamma_i$.

Definition 2.14. If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$, the simplex Δ_X is the set of Γ-trees equivariantly homeomorphic to X (component by component). The Euclidean topology and distance on Δ_X are defined by

$$d_\Delta(X, Y) = \sup_i d_{\Delta_{X_i}}(X_i, Y_i).$$

We notice that the simplicial structure of $\mathbb{P}\mathcal{O}(\Gamma)$ is not the product of the structures of $\mathbb{P}\mathcal{O}(\pi_1(\Gamma_i))$.

Remark 2.15. If $X \in \mathcal{O}(G)$, then $\mathcal{O}(X) = \mathcal{O}(G)$. In other words, $\mathcal{O}(G)$ is a particular case of $\mathcal{O}(\Gamma)$ whit Γ connected. In the following we will therefore develop the theory of $\mathcal{O}(\Gamma)$ and that of $\mathcal{O}(G)$ at once.

Definition 2.16 (Faces and closed simplices)**.** Let X be a G-graph (resp. a Γ-graph) and let $\Delta = \Delta_X$ be the corresponding open simplex. Let $F \subset X$ be a forest whose trees each contains at most one non-free vertex. The collapse of F in X produces a new G-graph (resp. Γ-graph), whence a simplex Δ_F. Such a simplex is called a *face* of Δ.

The *closed* simplex $\overline{\Delta}$ is defined by

$$\overline{\Delta} = \Delta \cup \{\text{all the faces of } \Delta\}.$$

2.4. **Simplicial bordification.** There are two natural topologies on $\mathcal{O}(G)$ (resp. $\mathcal{O}(\Gamma)$), the simplicial one and the equivariant Gromov topology, which are in general different. Here we will mainly use the simplicial topology. We notice that if Δ is an open simplex, the simplex $\overline{\Delta}$ is not the standard simplicial closure of Δ, because not all its simplicial faces are *faces* according to Definition 2.16. This is because some simplicial face of Δ are not in $\mathcal{O}(G)$ (resp. $\mathcal{O}(\Gamma)$) as defined. Such faces are somehow "at infinity" and describe limit points of sequences in $\mathcal{O}(G)$ (resp. $\mathcal{O}(\Gamma)$). We give now precise definitions to handle such limit points.

We will sometimes refer to the faces of Δ, as defined in Definition 2.16 as *finitary faces* of Δ.

Definition 2.17. Given an open simplex Δ in $\mathcal{O}(\Gamma)$, its *boundary at the finite* is the set of its proper faces:

$$\partial_{\mathcal{O}}\Delta = \partial_{\mathcal{O}}\overline{\Delta} = \overline{\Delta} \setminus \Delta.$$

Definition 2.18. A *core-graph* is a connected graph of groups whose leaves (univalent vertices) have non-trivial vertex-group. Given a graph X we define $\mathrm{core}(X)$ to be the maximal core sub-graph of X. (If the vertex groups are all trivial, so that X is simply a graph, then a core graph has no valence one vertices).

Note that $\mathrm{core}(X)$ is obtained by recursively cutting edges ending at leaves.

Let X be a Γ-graph and $\Delta = \Delta_X$. Let A be a proper subgraph of X having at least a component which is not a tree with at most one non-free vertex. Let Y be the graph of groups obtained by collapsing each component of A to a point (different components to different points). Then, $Y \in \mathcal{O}(X/A)$. The corresponding simplex Δ_Y is a simplicial face of Δ_X obtained by setting to zero the edge-lengths of A.

Definition 2.19. A face Δ_Y obtained as just described is called a *face at infinity* of $\overline{\Delta}_X$. If in addition we have that all components of A are core-graphs, then we say that Δ_Y is a face at infinity of Δ_X.

We define the *boundaries at infinity* by

$$\partial_\infty \Delta = \{\text{faces at infinity of } \Delta\}$$
$$\partial_\infty \overline{\Delta} = \{\text{faces at infinity of } \overline{\Delta}\},$$

and the *closure at infinity by*

$$\overline{\Delta}^\infty = \overline{\Delta} \cup \partial_\infty \overline{\Delta}.$$

If we denote by $\partial\Delta$ the simplicial boundary of Δ, we have

$$\partial\Delta = \partial_\infty \overline{\Delta} \cup \partial_{\mathcal{O}} \overline{\Delta}$$

and

$$\partial_\infty \overline{\Delta} = \bigcup_{F=\text{face of } \Delta} \partial_\infty F$$

(where the union is over all faces of Δ, Δ included.) Moreover, the simplicial closure of Δ is just $\overline{\Delta}^\infty$.

Definition 2.20. We define the boundary at infinity and the simplicial bordification of $\mathcal{O}(\Gamma)$ as

$$\partial_\infty \mathcal{O}(\Gamma) = \bigcup_{\Delta \text{ simplex}} \partial_\infty \Delta \qquad \text{and} \qquad \overline{\mathcal{O}(\Gamma)} = \overline{\mathcal{O}(\Gamma)}^\infty = \mathcal{O}(\Gamma) \cup \partial_\infty \mathcal{O}(\Gamma).$$

2.5. **Horospheres and regeneration.**

Definition 2.21. Given $X \in \partial_\infty \mathcal{O}(\Gamma)$, the horosphere $\mathrm{Hor}(\Delta_X)$ of Δ_X in $\mathcal{O}(\Gamma)$ is the union of simplices Δ such that $X \in \partial_\infty \Delta$. If $X \in \mathcal{O}(\Gamma)$ we set $\mathrm{Hor}(\Delta_X) = \Delta_X$.

The horosphere $\mathrm{Hor}(X)$ of X in $\mathcal{O}(\Gamma)$ is the set formed by points $Y \in \mathrm{Hor}(\Delta_X)$ such that $L_Y(e) = L_X(e)$ for any edge e of X. (In particular, if $X \in \mathcal{O}(\Gamma)$ we have $\mathrm{Hor}(X) = X$.)

Thus, Y is in the horosphere of X if X is obtained from Y by collapsing a proper family of core sub-graphs. On the other hand, $\mathrm{Hor}(X)$ can be regenerated from X as follows.

Suppose $X \in \partial_\infty \mathcal{O}(\Gamma)$. Thus there is a Γ-graph Y and a sub-graph $A = \sqcup_i A_i \subset Y$ whose components A_i are core-graphs, and such that $X = Y/A$. Let v_i be the non-free vertex of X corresponding to A_i. In order to recover a generic point $Z \in \mathrm{Hor}(X)$, we need to replace each v_i with an element $V_i \in \mathcal{O}(A_i)$. Moreover, in order to define the

marking on Z, we need to know where to attach to V_i the edges of X incident to v_i, and this choice has to be done in the universal covers $\widetilde{V}_i$. No more is needed. Therefore, if k_i denote the valence of the vertex v_i in X, we have

$$\mathrm{Hor}(X) = \Pi_i \mathcal{O}(A_i, k_i).$$

(Note that some k_i could be zero, e.g. if A_i is a connected component of Y.) There is a natural projection $\mathrm{Hor}(X) \to \mathcal{O}(A)$ which forgets the marking. We will be mainly interested in cases when we collapse A uniformly, for that reason we will use the projection

$$\pi : \mathrm{Hor}(X) \to \mathbb{P}\mathcal{O}(A)$$

where $\mathrm{Hor}(X)$ is intended to be not projectivized.

Note that if $[P] \in \mathbb{P}\mathcal{O}(A)$, then $\pi^{-1}(P)$ is connected because it is just $\Pi_i(A_i^{k_i})$. Since $\mathcal{O}(A)$ is connected, then $\mathrm{Hor}(X)$ is connected.

Finally, we notice that a graph of groups X can be considered as a point at infinity of different spaces. If we need to specify in which space we work we write $\mathrm{Hor}_\Gamma(X)$ or $\mathrm{Hor}_G(X)$.

2.6. The groups $\mathrm{Aut}(\Gamma)$ and $\mathrm{Out}(\Gamma)$.

Definition 2.22. Let G be endowed with the splitting $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$. The group of automorphisms of G that preserve the set of conjugacy classes of the G_i's is denoted by $\mathrm{Aut}(G; \mathcal{G})$. We set $\mathrm{Out}(G; \mathcal{G}) = \mathrm{Aut}(G; \mathcal{G})/\mathrm{Inn}(G)$[2].

The group $\mathrm{Aut}(G, \mathcal{G})$ acts on $\mathcal{O}(G)$ by changing the marking (i.e. the action), and $\mathrm{Inn}(G)$ acts trivially. Hence $\mathrm{Out}(G; \mathcal{G})$ acts on $\mathcal{O}(G)$. If $X \in \mathcal{O}(G)$ and $\phi \in \mathrm{Out}(G; \mathcal{G})$ then ϕX is the same metric tree as X, but the action is $(g, x) \to \phi(g)x$. The action is simplicial and continuous w.r.t. both simplicial and equivariant Gromov topologies.

We now extend the definition of $\mathrm{Aut}(G, \mathcal{G})$ to the case of $\Gamma = \sqcup_i \Gamma_i$. Let $\mathfrak{S}_k$ denotes the group of permutations of k elements.

Definition 2.23. Let G and H be two isomorphic groups endowed with splitting $\mathcal{G} : G = G_1 * \ldots G_p * F_n$ and $\mathcal{H} : H = H_1 * \ldots H_p * F_n$. The set of isomorphisms from G to H that maps each G_i to a conjugate of one of the H_i is denoted by $\mathrm{Isom}(G, H)$. If we need to specify the splittings we write $\mathrm{Isom}(G, H; \mathcal{G}, \mathcal{H})$.

Definition 2.24. For $\Gamma = \sqcup_{i=1}^k \Gamma_i$ as in Notation 2.10, we set

$$\mathrm{Aut}(\Gamma) = \{\phi = (\sigma, \phi_1, \ldots, \phi_k) : \ \sigma \in \mathfrak{S}_k \text{ and } \phi_i \in \mathrm{Isom}(H_i, H_{\sigma_i})\}.$$

The composition of $\mathrm{Aut}(\Gamma)$ is component by component defined as follows. Given $\phi = (\sigma, \phi_1, \ldots, \phi_k)$ and $\psi = (\tau, \psi_1, \ldots, \psi_k)$ we have

$$\psi\phi = (\tau\sigma, \psi_{\sigma(1)}\phi_1, \ldots, \psi_{\sigma(k)}\phi_k)$$

[2]Clearly $\mathrm{Inn}(G) \subset \mathrm{Aut}(G; \mathcal{G})$.

Remark 2.25. Not all permutations appear. For instance, if the groups H_i are mutually not isomorphic, then the only possible σ is the identity.

Definition 2.26. We set:

$$\mathrm{Inn}(\Gamma) = \{(\sigma, \phi_1, \dots, \phi_k) \in \mathrm{Aut}(\Gamma) : \sigma = id, \phi_i \in \mathrm{Inn}(H_i)\}$$
$$\mathrm{Out}(\Gamma) = \mathrm{Aut}(\Gamma)/\mathrm{Inn}(\Gamma).$$

Example 2.27. If $X \in \mathcal{O}(G)$ and $f : X \to X$ is a homotopy equivalence which leaves invariant a subgraph A, then $f|_A$ induces and element of $\mathrm{Aut}(A)$, and its free homotopy class an element of $\mathrm{Out}(A)$.

The group $\mathrm{Out}(\Gamma)$ acts on $\mathcal{O}(\Gamma)$ as follows. If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$, then each X_i is an H_i-tree. If $(\sigma, \phi_1, \dots, \phi_k) \in \mathrm{Aut}(\Gamma)$ then $X_{\sigma(i)}$ becomes an H_i-tree via the pre-composition of $\phi_i : H_i \to H_{\sigma(i)}$ with the $H_{\sigma(i)}$-action. We denote such an H_i-tree by $\phi_i X_{\sigma(i)}$. With that notation we have $\phi(X_i, \dots, X_n) = (\phi_1 X_{\sigma(1)}, \dots, \phi_k X_{\sigma(k)})$. (We remark that despite the left-positional notation, this is a right-action.)

3. PL-maps, gate structures, and optimal maps.

In this section we describe the theory of maps between graphs (or trees) representing points in outer spaces. We will treat in parallel the "connected" case $\mathcal{O}(G)$ and the general case $\mathcal{O}(\Gamma)$, where G and Γ are as in Notation 2.10.

3.1. PL-maps.

Now we will mainly work with trees.

Definition 3.1 ($\mathcal{O}$-maps in $\mathcal{O}(G)$)**.** Let $X, Y \in \mathcal{O}(G)$. A map $f : X \to Y$ is called an $\mathcal{O}$-map if it is Lipschitz-continuous and G-equivariant. The Lipschitz constant of f is denoted by $\mathrm{Lip}(f)$.

We recall that we tacitly identify $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ with the labelled disjoint union $\sqcup_i X_i$. Hence, if $X, Y \in \mathcal{O}(\Gamma)$, a continuous map $f : X \to Y$ is a collection of continuous maps $f_i; X_i \to Y_j$ for some $j = j(i)$.

Definition 3.2 ($\mathcal{O}$-maps in $\mathcal{O}(\Gamma)$)**.** Let $X = (X_1, \dots, X_k)$ and $Y = (Y_1, \dots, Y_k)$ be two elements of $\mathcal{O}(\Gamma)$. A map $f = (f_1, \dots, f_k) : X \to Y$ is called an $\mathcal{O}$-map if for each i the map f_i is an $\mathcal{O}$-map from X_i to Y_i.

Definition 3.3 (PL-maps)**.** Let X, Y be two metric trees. A Lipschitz-continuous map $f : X \to Y$ is a PL-map if it has constant speed on edges, that is to say, for any edge e of X there is a non-negative number $\lambda_e(f)$ such that for any $a, b \in e$ we have $d_Y(f(a), f(b)) = \lambda_e(f) d_X(a, b)$. If $X, Y \in \mathcal{O}(G)$ then we require any PL-map to be an $\mathcal{O}$-map. A PL-map between elements of $\mathcal{O}(\Gamma)$ is an $\mathcal{O}$-map whose components are PL. (If X, Y are Γ-graph, we understand that $f : X \to Y$ is a PL-map if its lift to the universal covers is PL:)

Remark 3.4. $\mathcal{O}$-map always exists and the images of non-free vertices is determined a priori by equivariance (see [5]). For any $\mathcal{O}$-map f there is a unique PL-map, denoted by $\mathrm{PL}(f)$, that conincides with f on vertices. We have $\mathrm{Lip}(\mathrm{PL}(f)) \leq \mathrm{Lip}(f)$.

Definition 3.5 ($\lambda_{\max}$ and tension graph)**.** Let $f : X \to Y$ be a PL-map. We set

$$\lambda(f) = \lambda_{\max}(f) = \max_e \lambda_e(f) = \mathrm{Lip}(f).$$

We define the *tension graph* of f by

$$X_{\max}(f) = \{e \text{ edge of } X : \lambda_e(f) = \lambda_{\max}\}.$$

If there are no ambiguities we set $\lambda_{\max} = \lambda_{\max}(f)$ and $X_{\max} = X_{\max}(f)$.

Definition 3.6 (Stretching factors)**.** For $X, Y \in \mathcal{O}(\Gamma)$ we define

$$\Lambda(X, Y) = \min_{f:X\to Y \ \mathcal{O}\text{-map}} \mathrm{Lip}(f)$$

The theory of stretching factors is well-developed in the connected case (i.e. for $\mathcal{O}(G)$), but one can readily see that connectedness of trees plays no role, and the theory extends without modifications to the non-connected case. In fact, Λ is well-defined, (see [6, 5]for details) and it satisfies the multiplicative triangular inequality:

$$\Lambda(X, Z) \leq \Lambda(X, Y)\Lambda(Y, Z)$$

It can be used to define a non-symmetric metric $d_R(X, Y) = \log(\Lambda(X, Y))$ and its symmetrized version $d_R(X, Y) + d_R(Y, X)$ (see [6, 7, 5] for details) which induces the Gromov topology. The group $\mathrm{Out}(\Gamma)$ acts by isometries on $\mathcal{O}(\Gamma)$.

Moreover, there is an effective way to compute Λ, via the so-called "sausage-lemma" (see [6, Lemma 3.14],[7, Lemma 2.16] for the classical case, and [5, Theorem 9.10] for the case of trees with non-trivial vertex-groups). We briefly recall here how it works.

Let $X, Y \in \mathcal{O}(\Gamma)$ (now seen as graphs). Any non-elliptic element $\gamma \in \pi_1(\Gamma)$ (i.e. an element not in a vertex-group) is represented by an immersed loop γ_X in X and one γ_Y in Y. The loop γ_X (or its lift to $\widetilde{X}$) is usually called **axis** of γ in X (or in $\widetilde{X}$). They have lengths $L_X(\gamma_X)$ and $L_Y(\gamma_Y)$ that correspond to the minimal translation length of the element γ acting on X and Y. (So $L_X(\gamma_X) = L_X(\gamma)$ and $L_Y(\gamma_Y)) = L_Y(\gamma)$.) We can define the stretching factor $L_Y(\gamma)/L_X(\gamma)$. Then $\Lambda(X, Y)$ is the minimum of the stretching factors of all non-elliptic elements.

Theorem 3.7 (Saussage Lemma [5, Theorem 9.10])**.** *Let $X, Y, \in \mathcal{O}(\Gamma)$. The stretching factor $\Lambda(X, Y)$ is realized by an element γ whose axis γ_X has one of the following forms:*

- *Embedded simple loop O;*
- *embedded "infinity"-loop ∞;*

- *embedded barbel* $O\!\!-\!\!O$*;*
- *singly degenerate barbel* $\bullet\!\!-\!\!O$*;*
- *doubly degenerate barbel* $\bullet\!\!-\!\!\bullet$*.*

(the $\bullet$ *stands for a non-free vertex.) Such loops are usually named "candidates".*

Remark 3.8. The stretching factor $\Lambda(X, Y)$ is defined on $\mathcal{O}(\Gamma)$ and not in $\mathbb{P}\mathcal{O}(\Gamma)$. However, we will mainly interested in computing factors of type $\Lambda(X, \phi X)$ (for $\phi \in \mathrm{Out}(\Gamma)$) and that factor is scale invariant.

Definition 3.9 (Gate structures)**.** Let X be any graph. A *gate* structure on X is an equivalence relation on germs of edges at vertices of X. Equivalence classes of germs are called *gates*. A *train-track* structure on X is a gate structure having at least two gates at every vertex. A *turn* is a pair of germs of edges incident to the same vertex. A turn is *illegal* if the two germs are in the same gate, it is *legal* otherwise. An immersed path in X is legal if it has only legal turns.

If $X = (X_1, \ldots, X_k) \in \mathcal{O}(\Gamma)$ we require the equivalence relation to be H_i-invariant on each X_i.

Any PL-map induces a gate structure as follows.

Definition 3.10 (Gate structure induced by f)**.** Given $X, Y \in \mathcal{O}(\Gamma)$ and a PL-map $f : X \to Y$, the gate structure induced by f, denoted by

$$\sim_f$$

is defined by declaring equivalent two germs that have the same non-degenerate f-image.

Remark 3.11 (See [5])**.** Given $X, Y \in \mathcal{O}(\Gamma)$ and $f : X \to Y$ a PL-map. If v is a non-free vertex of X and e if an edge incident to v which is not collapsed by f, then e and ge are in different gates for any $id \neq g \in \mathrm{Stab}(v)$.

Definition 3.12 (Optimal maps)**.** Given $X, Y \in \mathcal{O}(\Gamma)$, a map $f : X \to Y$ is *weakly optimal* if it is PL and $\lambda(f) = \Lambda(X, Y)$.

A map $f : X \to Y$ is *optimal* if the restriction of the gate structure induced by f, to the tension graph, is a train track structure (in other words, if the vertices of $X_{\max}$ are at least two-gated in $X_{\max}$).

Remark 3.13. Optimal maps always exist and are weakly optimal. A map between two Γ-trees is weakly optimal if and only there is an periodic immersed legal line in the tension graph (i.e. a legal immersed loop in the quotient graph).

In general optimal maps are neither unique nor do they form a discrete set, even if $X_{\max} = X$, as the following example shows. (If $X_{\max} \neq X$ then one can use freedom given by the lengths of edges not in $X_{\max}$ to produce examples.)

Example 3.14 (A continuous family of optimal maps with $X_{\max} = X$)**.** Consider $G = F_2$. Let X be a graph with three edges e_1, e_2, e_3 and two free vertices P, Q, as in Figure 1. Set the length of e_2 to be 2, name x the length of e_1, and $1+\delta$ that of e_3. The parameters x, δ will be determined below. For any $t \in [0,1]$ consider the point P_t at distance $1+t$ from P along e_2, and the point Q_t at distance $1-t$ from P along e_3. P_t divides e_2 in oriented segments a_t, c_t. Q_t divides e_3 into b_t, d_t. Consider the PL-map $f : X \to X$ defined as in the Figure, sending P

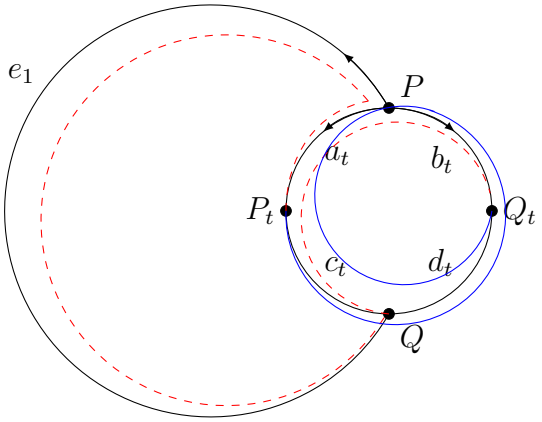

$e_2 = a_t c_t$
$e_3 = b_t d_t$
$f_t(e_1) = \bar{a}_t e_1 \bar{c}_t \bar{a}_t b_t$
$f_t(e_2) = c_t \bar{d}_t$
$f_t(e_3) = c_t \bar{d}_t \bar{b}_t a_t c_t \bar{d}_t$
$\text{Length}(a_t) = 1+t$
$\text{Length}(b_t) = 1-t$
$\text{Length}(c_t) = 1-t$
$\text{Length}(d_t) = \delta + t$
$\text{Length}(e_1) = x$
$\text{Length}(e_2) = 2$
$\text{Length}(e_3) = 1+\delta$

FIGURE 1. A continuous family of optimal maps with $X_{\max} = X$. The red dashed line is $f(e_1)$ and the blue line is $f(e_3)$ ($f(e_2)$ is not depicted).

to P_t and Q to Q_t. If we collapse e_3, and we homotop P_t to P along a, this corresponds to the automorphism $e_1 \mapsto e_1\overline{e_2}, e_2 \mapsto \overline{e_2}$.

A direct calculation shows that if we set $\delta = 1+2\sqrt{2}$ and $x = 2\sqrt{2}$, the map f_t is optimal for any t and all the three edges are stretched by the same amount, as follows.

The edges e_1 and e_2 are in different gates at P and e_1 and e_3 are in different gates at Q. In order to check that f_t is optimal it suffices to check that every edge is stretched the same.

$$\lambda_{e_1}(f_t) = \frac{x+4}{x} \qquad \lambda_{e_2}(f_t) = \frac{1+\delta}{2} \qquad \lambda_{e_3}(f_t) = \frac{4+2\delta}{1+\delta}.$$

In particular they do not depend on t. If we set $x = 2\sqrt{2}$ and $\delta = 1+2\sqrt{2}$ we get

$$\lambda_{e_1}(f_t) = \frac{2\sqrt{2}+4}{2\sqrt{2}} \qquad \lambda_{e_2}(f_t) = \frac{2+2\sqrt{2}}{2} \qquad \lambda_{e_3}(f_t) = \frac{6+4\sqrt{2}}{2+2\sqrt{2}}$$

and all of them are $1+\sqrt{2}$. □

However, given a PL-map, we can choose an optimal map which is in some sense the closest possible. Given two $\mathcal{O}$-maps $f, g : X \to Y$ we define

$$d_\infty(f,g) = \max_{x\in X} d_Y(f(x), g(x)).$$

For $X \in \mathcal{O}(G)$ we define its (co-)volume $\mathrm{vol}(X)$ as the sum of lengths of edges in $G\backslash X$. If $X = (X_1, \dots, X_k) \in \mathcal{O}(\Gamma)$ we set $\mathrm{vol}(X) = \sum_i \mathrm{vol}(X_i)$.

Theorem 3.15 (Optimization)**.** *Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \to Y$ be a* PL*-map. There is a map*[3] $\mathrm{weakopt}(f) : X \to Y$ *which is weakly optimal and such that*

$$d_\infty(f, \mathrm{weakopt}(f)) \leq \mathrm{vol}(X)(\lambda(f) - \Lambda(X,Y))$$

Moreover, for any $\varepsilon > 0$ there is an optimal map $g : X \to Y$ such that $d_\infty(g, \mathrm{weakopt}(f)) < \varepsilon$.

Proof. By arguing component by component, we may assume without loss of generality that Γ is connected, hence that we are in the case $X, Y \in \mathcal{O}(G)$. For this proof it will be convenient to work with graphs rather than with trees (so $X = G\backslash\widetilde{X}$, whit $\widetilde{X}$ a G-tree). By Remark 3.11 a non-free vertex will never be considered one-gated (because it is never one-gated in $\widetilde{X}$).

Let us concentrate on the first claim.

Let $\lambda = \Lambda(X,Y)$. Since PL-maps are uniquely determined by their value on vertices, we need only to define $\mathrm{weakopt}(f)$ (and g) on vertices of X. By Remark 3.4 the image of non-free vertices is fixed. We define PL-maps f_t for $t \in [0, \lambda_f - \lambda]$ by moving the images of all one-gated vertices of $X_{\max}(f_t)$, in the direction given by the gate, so that

$$\frac{d}{dt}\lambda(f_t) = -1.$$

Let us be more precise on this point. We define a flow which is piecewise linear, depending on the geometry of the tension graph at time t. The key remark to have in mind is that if an edge is not in $X_{\max}(f)$, then it remains in the complement of the tension graph for small perturbations of f. Therefore, we can restrict our attention to the tension graph.

Suppose we are at time t. We recursively define sets of vertices and edges as follows:

- V_0 is the set vertices of $X_{\max}(f_t)$ which are one-gated in $X_{\max}(f_t)$;
- E_0 is the set of edges of $X_{\max}(f_t)$ incident to vertices in V_0. We agree that such edges contain the vertices in V_0 but not others. (If and edge has both vertices in V_0 then it contains both, otherwise it contains only one of its vertices.)

[3]We describe an algorithm to fine the map $\mathrm{weakopt}(f)$, but the algorithm will depend on some choice, hence the map $\mathrm{weakopt}(f)$ may be not unique in general.

Having defined $V_0, \dots, V_i$ and $E_0, \dots, E_i$, we define V_{i+1} and E_{i+1} as follows:

- V_{i+1} is the set of one-gated vertices of $X_{\max}(f_t) \setminus \cup_{j=0}^{i} E_j$;
- E_{i+1} is the set of edges of $X_{\max}(f_t) \setminus \cup_{i=0}^{i} E_i$ incident to vertices in V_{i+1}. (As above such edges contain vertices in V_{i+1} but not others.)

We notice that since X is a finite G-graph, we have only finitely many sets V_i, say $V_0, \dots, V_k$.

Lemma 3.16. *If f_t is not weakly optimal, then $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$ is a (possibly empty) collection of vertices, that we name terminal vertices.*

Proof. Note that no vertex in $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$ can be one-gated, hence any vertex in $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$ is either isolated or has at least two gates in $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$. Thus if there is an edge e in $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$, the component of $X_{\max}(f_t) \setminus \cup_{i=0}^{k} E_i$ containing e must also contain an immersed legal loop and so f_t is weakly optimal. □

By convention we denote the set of terminal vertices by V_∞.

Remark 3.17. Any $e \in E_i$ has by definition at least one endpoint in V_i, and the other endpoint is in some V_j with $j \geq i$.

Our flow is defined by moving the images $f_t(v)$ of vertices in $X_{\max}(f_t)$. We need to define a direction and a speed $s(v) \geq 0$ for any $f_t(v)$.

For $i < \infty$ each vertex in V_i has a preferred gate: the one that survives in $X_{\max}(f_t) \setminus \cup_{j=0}^{i-1} E_j$. That gate gives us the direction in which we move $f_t(v)$.

Thus a vertex in V_0 is one-gated, and hence we define the flow so as to reduce the Lipschitz constant for every edge in E_0 (shrinking the image of each E_0 edge). Similarly, every vertex in V_1 is one gated in $X_{\max}(f_t) \setminus E_0$, so we define the flow to reduce the Lipschitz constants of edges in E_1 and so on.

Now we define the speeds.

Lemma 3.18. *There exists speeds $s(v) \geq 0$ such that if we move the images of any v at speed $s(v)$ in the direction of its preferred gate, then for any edge $e \in X_{\max}(t)$*

$$\frac{d}{dt}\lambda_e(f_t) \leq -1.$$

Moreover, for any i, and for any $v \in V_i$, either $s(v) = 0$ or there is an edge $e \in E_i$ incident at v such that

$$\frac{d}{dt}\lambda_e(f_t) = -1.$$

Proof. We start by giving a total order of the vertices of $X_{\max}(f_t)$ in such a way that vertices in V_i are bigger than those in V_j whenever $i > j$. We define the speeds recursively.

The speed of terminal vertices is set to zero. Let v be a vertex of $X_{\max}(f_t)$ and suppose that we already defined the speed $s(w)$ for all $w > v$.

The vertex v belongs to some set V_i. For any edge $e \in E_i$ emanating from v let u_e be the other endpoint of of e, and define a sign $\sigma_e(u_e) = \pm 1$ as follows: $\sigma_e(u_e) = -1$ if the germ of e at u_e is in the preferred gate of u_e, and $\sigma_e(u_e) = 1$ otherwise. (So, for example, $\sigma_e(u_e) = 1$ if u_e is terminal, and $\sigma_e(u_e) = -1$ if $v = u_e$, or if $u_e \in V_i$.)

With this notation, if we move $f(v)$ and $f(u_e)$ in the direction given by their gates at speeds $s(v)$ and ν respectively, then the derivative of $\lambda_e(f_t)$ is given by

$$-\frac{s(v) - \sigma_e(u_e)\nu}{L_X(e)}$$

If $u_e > v$ we already defined its speed. We set

$$s(v) = \max\{0, \max_{u_e > v}\{L_X(e) + \sigma_e(u_e)s(u_e)\}, \max_{u_e = v} \frac{L_X(e)}{2}\}$$

where the maxima are taken over all edges $e \in E_i$ emanating from v. Note that there may exist some such edge with $u_e < v$. By Remark 3.17 in this case $u_e \in V_i$ (same i as v), $\sigma_e(u_e) = -1$ and the derivative of λ_e will be settled later, when defining the speed of u_e.

With the speeds defined in that way we are sure that for any edge e we have $d/dt\lambda_e(f_t) \leq -1$ and, if $s(v) \neq 0$, then the edges that realize the above maximum satisfy $d/dt\lambda_e(f_t) = -1$. □

Locally, when we start moving, the tension graph may lose some edges. However, the above lemma assures that any vertex v with $s(v) \neq 0$ is incident to an edge e which is maximally stretched and $d/dt\lambda_e = -1$. Hence such an edge remains in the tension graph when we start moving. Since $d/dt\lambda_e \leq -1$ for any edge, it follows that when we start moving, the tension graph stabilizes. So our flow is well defined in $[t, t + \epsilon]$ for some $\epsilon > 0$. If at a time $t_1 > t$ some edge that was not previously in $X_{\max}(f_t)$ becomes maximally stretched we recompute speeds and we start again. A priori we may have to recompute speeds infinitely many times $t < t_1 < t_2 < \dots$ but the control on $d/dt\lambda(f_t)$ assures that $\sup t_i = T \leq \lambda_f - \lambda$. Since speeds are bounded the flow has a limit for $t \to T$ an then we can restart from T. Therefore the set of times $s \in [0, \lambda_f - \lambda]$ for which the flow is well-defined for $t \in [0, s]$ is closed and open and thus is the whole $[0, \lambda_f - \lambda]$.

With these speeds, we have $d/dt(\lambda(f_t)) = -1$. Therefore for $t = \lambda(f) - \lambda$, and not before, we have $\lambda(f_t) = \lambda$ hence f_t is weakly optimal. We define

$$\mathrm{weakopt}(f) = f_{\lambda(f)-\lambda}.$$

Now we estimate $d_\infty(f, f_t)$. The d_∞-distance between PL-maps is bounded by the d_∞-distance of their restriction to vertices.

We first estimate the speed at which the images of vertices move. Let S be the maximum speed of vertices, i.e. $S = \max_v |s(v)|$. Let v be a fastest vertex. Since it moves, it belongs to V_s for some $s < \infty$. Let $v = v_1, v_2 \dots, v_m$ be a maximal sequence of vertices such that:

(1) $s(v_i) > 0$ for $i < m$;
(2) there is an edge e_i between v_i and v_{i+1} such that $e_i \in E_a$ if $v_i \in V_a$;
(3) $\sigma_{e_i}(v_{i+1}) = 1$ for $i + 1 < m$;
(4) $d/dt(\lambda_{e_i}(f_t)) = -1$.

By the above lemma, we have that either $s(v_m) = 0$ or $\sigma_{e_{m-1}}(v_m) = -1$. Moreover, by $(2) - (3)$ and Remark 3.17 we have that $v_i < v_{i+1}$ and therefore the edges e_i are all distinct.

Let γ be the path obtained by concatenating the e_i's. By $(2) - (3)$, γ is a legal path in the tension graph. So let

$$L = \sum_i L_X(e_i) = L_X(\gamma) \qquad L_t = \sum_i L_Y(f_t(e_i)) = L_Y(f_t(\gamma)).$$

Since the e_i's are in the tension graph and by condition (4) we have

$$L_t = \lambda(f_t)L \qquad \frac{d}{dt}L_t = -L$$

On the other hand $-\frac{d}{dt}L_t \geq S$ because by (3) the contributions of the speeds of v_i does not count for $i = 2, \dots, m-1$ and $f(v_m)$ either stay or moves towards $f(v_1)$. It follows that

$$S \leq L \leq \operatorname{vol}(X).$$

It follows that for any vertex v we have

$$d_Y(f(v), f_t(v)) \leq \int_0^t \left| \frac{d}{ds} f_s(w) \right| ds \leq \int_0^t S = tS \leq t \operatorname{vol}(X)$$

hence

$$d_\infty(\text{weakopt}(f), f)) = d_\infty(f_{\lambda(f)-\lambda}, f) \leq (\lambda(f) - \lambda) \operatorname{vol}(X).$$

We prove now the last claim. If weakopt(f) is optimal then we are done. Otherwise, there is some one-gated vertex in $X_{\max}$. We start moving the one-gated vertices as described above, for an arbitrarily small amount. Let g be the map obtained, clearly we can make $d_\infty(g, \text{weakopt}(f))$ arbitrarily small. Since weakopt(f) is optimal, we must have $\lambda(g) = \lambda(\text{weakopt}(f))$. It follows that there is a core sub graph of $X_{\max}$ which survives the moving. In particular, every vertex of $X_{\max}(g)$ is at least two-gated, hence g is optimal. □

Definition 3.19. We denote by opt(f) any optimal map obtained from weakopt(f) as described in the proof of Theorem 3.15.

Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \to Y$ be an optimal map. Let v be a vertex of X having an f-illegal turn $\tau = (e_1, e_2)$. Since $f(e_1)$ and $f(e_2)$ share an initial segment, we can identify an initial segment of e_1 and e_2. We obtain a new element $X' \in \mathcal{O}(\Gamma)$, with an induced optimal map, still denoted by f, $f : X' \to Y$. This is a particular case of Stallings folding ([17]). We refer to [5] for further details.

Definition 3.20. We call such an operation **simple fold directed by** f.

We finish this section by proving the existence of optimal maps with an additional property, that will be used in the sequel.

Definition 3.21. Let $X, Y \in \mathcal{O}(\Gamma)$. An optimal map $f : X \to Y$ is *minimal* if its tension graph consists of the union of axes of maximally stretched elements it contains. In other words, if any edge $e \in X_{\max}$ is contained in the axis of some element in $\pi_1(X_{\max})$ which is maximally stretched by f.

Note that not all optimal maps are minimal, as the following shows.

Example 3.22. Let X be the graph consisting of two barbels joined by an edge, as in Figure 2. All edges have length one except the two lower loops that have length two.

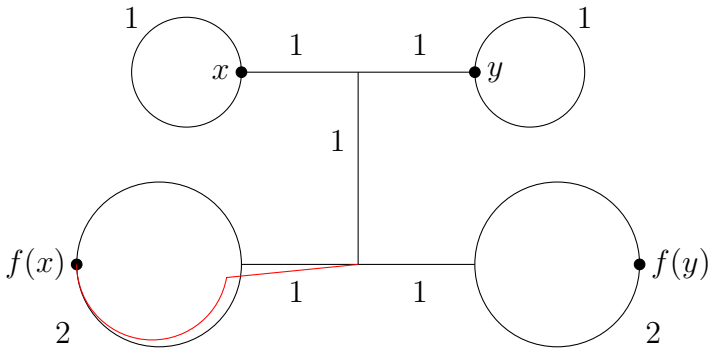

FIGURE 2. A non-minimal optimal map. The dots $f(x)$ and $f(y)$ are not vertices, all other crossings are. The red line is the image of the left "bar-edge" of the top barbell.

Let $f : X \to X$ be the PL-map that exchanges the the top and bottom barbells (preserving left and right) and maps x to the middle point of the lower left loop, and y to the middle point of the lower right loop (see the figure).

The restriction of f to the lower barbell is 1-Lipschitz (each loop is shrunk and the bar is the same length as its image), while the stretching factor of all top edges is two. Hence the tension graph $X_{\max}$ is the top barbel. The map is optimal because all vertices of $X_{\max}$ are two

gated, but the "bar-edges" of the top barbel are not in the axis of any maximally stretched loop. This is because the only legal loops in $X_{\max}$ are the two lateral loops of the barbell. Clearly this map can be homotoped to a map with smaller tension graph. As the next theorem shows this is always the case for non-minimal optimal maps. □

Theorem 3.23. *Let $X, Y \in \mathcal{O}(\Gamma)$ and let $f : X \to Y$ be an optimal map. If f locally minimizes the tension graph amongst all optimal maps $X \to Y$, then f is minimal. Moreover, given $g : X \to Y$ optimal, for any $\varepsilon > 0$ there is a minimal optimal map $f : X \to Y$ with $d_\infty(g, f) < \varepsilon$.*

Proof. The first claim clearly implies the second, because the tension graph is combinatorially finite, hence the set of possible tension graphs is finite and we can always locally minimize it.

We will prove the contrapositive, that if f is not minimal then we can decrease the tension graph by perturbations as small as we want. The spirit is similar to that of the proof of Theorem 3.15. Again we will work with graphs. At the level of graphs, the non-minimality of f translates to the fact that there is an edge α in the tension graph which is not part of any legal loop in $X_{\max}$.

Let x be the terminal vertex of the edge α. We say that a path starting at x is α-legal, if it is a legal path in the tension graph, whose initial edge, e, is not in the same gate as $\overline{\alpha}$. We say a loop at x is α-legal if, considered as paths, both the loop and its inverse are α-legal.

If the terminal vertex of α admits an α-legal loop and the initial point of α also admits an $\overline{\alpha}$-legal loop, then we can form the concatenation of these loops with α to get a legal loop in the tension graph crossing α and contradicting our hypothesis. Hence we may assume that the endpoint (rather than the initial point) of α is a vertex, x, which admits no α-legal loops.

We will show that it is possible to move the image of x a small amount (and possibly some other vertices) so that we obtain an optimal map with smaller tension graph. Let ε small enough so that if an edge is not in $X_{\max}$, than it remains outside the tension graph for any perturbation of f by less than ε.

From now on, we restrict ourselves to the tension graph. We say that a vertex v is legally seen from x if there is an α-legal path γ from x to v. Note that in this case v is free. Indeed, otherwise the path γ followed by its inverse is in fact an α-legal loop (it has a legal lift to $\widetilde{X}$ defined by using the action of the stabilizer of v). Since v is free, we can move $f(v)$. Also observe that the initial point of α is not α-legally seen from x, since otherwise we would get a legal loop in the tension graph containing α.

We want to chose a direction to move the images of vertices α-legally seen from x. First, the direction we choose for $f(x)$ is given by the gate

of α. That is, we move $f(x)$ so as to reduced the length of α. For any vertex, v, α-legally seen from x, via a path γ, we move $f(v)$ backwards via the last gate of γ. That is, we move $f(v)$ so as to retrace γ. Note that this direction depends only on v and not on the choice of γ. This is because, were there to be another α-legal path from x to v, γ', then the concatenation $\gamma\overline{\gamma'}$ would define an α-legal loop at x unless the terminal edges of γ, γ' lie in the same gate. Hence directions are well defined.

We move the images of all vertices by ε in the direction given above. Consider an edge, β (not equal to α or its inverse) in the tension graph. If neither vertex of β is α-legally seen from x, then the image of β is unchanged and it remains in the tension graph. Otherwise, suppose that the initial vertex of β is α-legally seen from x, via a path γ, whose terminal edge is $\overline{\eta}$. If η and β are in different gates, then the terminal vertex of β is also α-legally seen from x and both vertices are moved the same amount, such that the length of the image of β remains unchanged. If, conversely, η and β are in the same gate then either the length of the image of β is reduced (if the terminal vertex is not α-legally seen) or it remains unchanged (if it is). Moreover, since the initial vertex of α is not α-legally seen, the length of the image of α must strictly decrease. In particular, α itself is no longer in the tension graph.

On the other hand, since the tension graph has no one-gated vertices, there is at least one α-legal path emanating from x, an so some part of the tension graph survives. Since f is optimal, our assumption on ε implies that the new map is optimal and it has a tension graph strictly smaller than f. □

4. Displacement function and train track maps for automorphisms

This section is devoted to the study of train track maps from a metric point of view. The spirit is that of [3, 5]. For the rest of the section we fix G and $\Gamma = \sqcup_i \Gamma_i$ as in Notation 2.10. We recall that the study of $\mathcal{O}(G)$ is a particular case of $\mathcal{O}(\Gamma)$ when Γ is connected.

We recall the main facts proved in [5] for irreducible elements of $\mathrm{Aut}(G)$, and we generalize such facts to the case of $\mathrm{Aut}(\Gamma)$, including reducible automorphisms. Connectedness does not really play a crucial role, and most of the proves of [5] work exactly in the same way. The key here is the passage from irreducible to reducible automorphisms.

For the rest of the section, if not specified otherwise, $\phi = (\sigma, \phi_1, \dots, \phi_k)$ will be an element of $\mathrm{Aut}(\Gamma)$. By abuse of notation, we will make no distinction between ϕ and the element of $\mathrm{Out}(\Gamma)$ it represents. We let act the symmetric group $\mathfrak{t}_k$ on $\mathcal{O}(\Gamma)$ by permuting the components: $\sigma(X_1, \dots, X_k) = (X_{\sigma(1)}, \dots, X_{\sigma(k)})$.

As usual, if there is no ambiguity we will make no distinction between Γ-graphs and Γ-trees. If necessary we well use $\widetilde{X}$ to refer to the Γ-tree

corresponding to a Γ-graph X. (So $\widetilde{X}$ will be the universal covering of X:)

Definition 4.1. Let $X \in \mathcal{O}(\Gamma)$. We say that a (PL) map $f = (f_1, \dots, f_k) : X \to X$ *represents* ϕ if f_i maps X_i to $X_{\sigma(i)}$, and $\sigma \circ f : X \to \phi X$ is an $\mathcal{O}$-map (resp. PL-map). We say that f is optimal if $\sigma \circ f$ is optimal.

If $\sigma \circ f : X \to \phi X$ is an optimal map, then any fold directed by f gives a new point X' as well as a new map, still denoted by f, such that f represents ϕ and $\sigma \circ f : X' \to \phi X'$ is an optimal map. (This follows from Theorem 3.15, see [5] for more details.)

Definition 4.2 (Displacements)**.** For any automorphism $\phi \in \mathrm{Out}(\Gamma)$ we define the function

$$\lambda_\phi : \mathcal{O}(\Gamma) \to \mathbb{R} \qquad \lambda_\phi(X) = \Lambda(X, \phi X)$$

If Δ is a simplex of $\mathcal{O}(\Gamma)$ we define

$$\lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$$

If there is no ambiguity we write simply λ instead of λ_ϕ. Finally, we set

$$\lambda(\phi) = \inf_{X \in \mathcal{O}(\Gamma)} \lambda_\phi(X)$$

Definition 4.3 (Minimally displaced points)**.** For any automorphism ϕ we define sets:

$$\mathrm{Min}(\phi) = \{X \in \mathcal{O}(\Gamma) : \lambda(X) = \lambda(\phi)\}$$

$$\mathrm{LocMin}(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists U \ni X \text{ open s.t. } \forall Y \in U \ \lambda(X) \leq \lambda(Y)\}$$

Remark 4.4. A fold directed by a weakly optimal map does not increase λ. In particular, $\mathrm{Min}(\phi)$ is invariant by folds directed by weakly optimal maps.

Definition 4.5 (Reducibility)**.** An automorphism ϕ is called *reducible* if there is a Γ-graph $X \in \mathcal{O}(\Gamma)$ and $f : X \to X$ representing ϕ having a proper f-invariant subgraph $Y \subset X$ such that at least a component of Y is not a tree with at most one non-free vertex.

Equivalently, ϕ is reducible if the above Γ contains the axis of a hyperbolic element.

ϕ is *irreducible* if it is not reducible.

Remark 4.6. In the connected case, if $G = F_n$ then this definition coincides with the usual definition of irreducibility. For irreducible automorphisms we have $\mathrm{Min}(\phi) \neq \emptyset$. (See [5] for more details.)

Definition 4.7 (Thin and thick reducible automorphisms)**.** A reducible automorphism ϕ is called *thick* if $\mathrm{Min}(\phi) \neq \emptyset$, and it is called *thin* otherwise.

Definition 4.8 (Thin and thick simplices)**.** Given a (reducible) automorphism ϕ, a simplex Δ of $\mathcal{O}(\Gamma)$ is called *ϕ-thick* (or simply *thick* for for short) if

$$\inf_{\Delta} \lambda_\phi \text{ is realized at a point of } \overline{\Delta}.$$

Otherwise Δ is called *ϕ-thin*. (Recall that $\overline{\Delta}$ means the finitary closure of Δ.)

Remark 4.9. If ϕ is irreducible, then any simplex is ϕ-thick. (See for instance [5, Section 8]. See also Proposition 5.6 below.) In [3, 5] irreducible and thick reducible automorphisms are called hyperbolic.

Definition 4.10 (Train track between trees)**.** Let $\sim$ be a gate structure on a (not necessarily connected) tree X. A map $f : X \to X$ is a **train track map w.r.t.** $\sim$ if

(1) $\sim$ is a train track structure (i.e. vertices have at least two gates);
(2) f maps edges to legal paths (in particular, f does not collapse edges);
(3) for any vertex v, if $f(v)$ is a vertex, then f maps inequivalent germs at v to inequivalent germs at $f(v)$.

We already defined the gate structure $\sim_f$ induced by a PL-map (Definition 3.10).

Definition 4.11. Let X be a (not necessarily connected) tree, and let $f : X \to X$ be a maps whose components are PL. We define the gate structure $\langle \sim_{f^k} \rangle$ as the equivalence relation on germs generated by all $\sim_{f^k},\ k \in \mathbb{N}$.

Lemma 4.12. *Let $\phi \in \mathrm{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$ and $\sim$ be a gate structure on X. Let $f : X \to X$ be a PL-map representing ϕ. If $f : X \to X$ is a train track map w.r.t. $\sim$, then $\sim \supseteq \langle \sim_{f^k} \rangle$. In particular if f is a train track map w.r.t. some $\sim$ then it is a train track map w.r.t $\langle \sim_{f^k} \rangle$.*

See [5, Section 8] for a proof (where it is proved in the connected case, but connectedness plays no role).

Now we give a definition of train track map representing an automorphism. Our definition is given at once for both reducible and irreducible automorphisms, and in the irreducible case is the standard one. This definition well-behaves with respect to the displacement function in the reducible case.

Definition 4.13 (Optimal train track maps for automorphisms)**.** Let $\phi \in \mathrm{Out}(\Gamma)$. Let X be a Γ-graph in $\mathcal{O}(\Gamma)$ and let $f : X \to X$ be a PL-map representing ϕ. Then we say that f is a

- **strict train track map** if there is an f-invariant sub-graph $Y \subseteq X_{\max}(f)$ such that $f|_Y$ is a train track map w.r.t. $\sim f$.

- **train track map** if there is an f-invariant sub-graph $Y \subseteq X_{\max}(f)$ such that $f|_Y$ is a train track map w.r.t. $\langle \sim_{f^k} \rangle$.

Here some remarks are needed. The theory of train tracks maps, introduced in [4], does not have a completely standard terminology, especially for reducible automorphisms. We want to describe the main properties of train track maps, comparing topological and metric viewpoints. Usually, *topological* train track maps are defined without requiring that the f-invariant sub-graph is in the tension graph.[4]

In the case ϕ is irreducible there is no much difference. Indeed if $f : X \to X$ is a topological train track map representing ϕ, then one can rescale the edge-lengths of X so that f is a train track map for Definition 4.13. And the same holds true if f has no proper invariant sub-graphs. This is because train track maps does not collapse edges, hence edge-lengths can be adjusted so that every edge is stretched the same. In particular, the following two results are proved in [5] for irreducible automorphisms and Γ connected. The proves for generic automorphisms are basically the same (details are left to the reader).

Lemma 4.14. *Let $\phi \in \mathrm{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$, and $f : \widetilde{X} \to \widetilde{X}$ be a* PL*-map representing ϕ. Then f is train track if and only if there is an immersed periodic line L in $\widetilde{X}_{\max}$ such that $f^k(L) \subseteq \widetilde{X}_{\max}$ and $f^k|_L$ is injective for all $k \in \mathbb{N}$. In particular if f is train track then*

1. *f^k is train track;*
2. *$\mathrm{Lip}(f) = \Lambda(X, \phi X)$ (hence f is weakly optimal);*
3. *$\mathrm{Lip}(f)^k = \mathrm{Lip}(f^k) = \Lambda(X, \phi^k X)$.*

Corollary 4.15. *Let $\phi \in \mathrm{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$, and $f : \widetilde{X} \to \widetilde{X}$ be a map representing ϕ. Suppose that there is an embedded periodic line L in $\widetilde{X}$ such that $f^k|_L$ is injective for all $k \in \mathbb{N}$. Suppose moreover that $\cup_k f^k(L) = \widetilde{X}$. Then there is X' obtained by rescaling edge-lengths of X so that* $\mathrm{PL}(f) : \widetilde{X}' \to \widetilde{X}'$ *is a train track map.*

In general, if $\cup_k f^k(L)$ is just an f-invariant subtree Y of $\widetilde{X}$, we can adjust edge lengths so that every edge of Y is stretched the same, but we cannot guarantee a priori that $Y \subset X_{\max}$. However, the *interesting* case is when $\cup_k f^k(L) = \widetilde{X}$.

Definition 4.16 (Train track sets)**.** For any automorphism $\phi \in \mathrm{Aut}(\Gamma)$ we define sets:

$$\mathrm{TT}(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists f : X \to X \text{ train track}\}$$

$$\mathrm{TT}_0(\phi) = \{X \in \mathcal{O}(\Gamma) : \exists f : X \to X \text{ strict train track}\}$$

[4]Our present definition of train track map coincides with the notion of *optimal* train track map given in [5] for irreducible automorphisms in the connected case.

If we need to specify the map we write $(X, f) \in \mathrm{TT}(\phi)$ or $(X, f) \in \mathrm{TT}_0(\phi)$.[5]

Theorem 4.17. *Let $\phi = (\sigma, \phi_1, \dots, \phi_k) \in \mathrm{Aut}(\Gamma)$. Then*

$$\overline{\mathrm{TT}_0(\phi)} = \mathrm{TT}(\phi) = \mathrm{Min}(\phi) = \mathrm{LocMin}(\phi)$$

where the closure is made with respect to the simplicial topology.

Proof. If ϕ is irreducible and Γ connected, the proof is given in [5] and goes trough the following steps:

(1) $\mathrm{TT}_0(\phi) \subset \mathrm{TT}(\phi) \subseteq \mathrm{Min}(\phi)$.
(2) If X locally minimizes λ_ϕ in Δ_X then $X_{\max} = X$.
(3) $\mathrm{TT}_0(\phi)$ is dense in $\mathrm{LocMin}(\phi)$.
(4) $\mathrm{TT}(\phi)$ is closed.

We now adapt the proof so that it works also for ϕ reducible and general Γ. Clearly $\mathrm{Min}(\phi) \subseteq \mathrm{LocMin}(\phi)$. By Lemma 4.12 $\mathrm{TT}_0(\phi) \subseteq \mathrm{TT}(\phi)$. We see now that $\mathrm{TT}(\phi) \subseteq \mathrm{Min}(\phi)$. If $X \in \mathrm{TT}(\phi)$ and $\lambda(X) > \lambda(\phi)$ then there is $Y \in \mathcal{O}(\Gamma)$ such that $\lambda(Y) < \lambda(X)$. By Lemma 4.14 $\Lambda(X, \phi^k X) = \lambda(X)^k$ but then

$$\lambda(X)^k = \Lambda(X, \phi^k X) \leq \Lambda(X, Y)\Lambda(Y, \phi^k Y)\Lambda(\phi^k Y, \phi^k X)$$
$$= \Lambda(X, Y)\Lambda(Y, \phi^k Y)\Lambda(\phi^k Y, \phi^k X) \leq \Lambda(X, Y)\Lambda(Y, X)\lambda(Y)^k$$

thus $(\frac{\lambda(X)}{\lambda(Y)})^k$ is bounded for any k, which is impossible if $\frac{\lambda(X)}{\lambda(Y)} > 1$.

Thus we have

$$\mathrm{TT}_0(\phi) \subseteq \mathrm{TT}(\phi) \subseteq \mathrm{Min}(\phi) \subseteq \mathrm{LocMin}(\phi).$$

Lemma 4.18. *Suppose (X, f) locally minimizes λ in Δ_X. Then there is $Y \subseteq X_{\max}$ which is f-invariant.*

Proof. For every open neighbourhood U of X we choose a point X^U such that

- it locally minimizes λ (a priori X^U can be X)
- it locally minimizes the tension graph with respect to the inclusion.

We still denote by f the optimal map $f : X^U \to X^U$ obtained by optimizing f w.r.t. the metric of X^U (see Theorem 3.15). If $f(X^U_{\max})$ contains an edge e which is not in the tension graph, then by shrinking a little such edge, either we reduce λ — which is impossible — or we reduce the tension graph — which is impossible too —. Thus $X^U_{\max}$ is f-invariant. By choosing a family of nested neighbourhoods U_i we provide a sequence $X^{U_i} \to X$ having an invariant subgraph in the tension graph. At the limit we get an invariant subgraph of the tension graph of X. □

[5]We remark that, since in the irreducible case our present definition of train track map corresponds to that of *optimal* train track map of [5], the two definitions of TT and TT_0 coincide with those given in [5].

Lemma 4.19. $\mathrm{LocMin}(\phi) \subseteq \overline{\mathrm{TT}_0(\phi)}$. *More precisely, let $X \in \mathcal{O}(\Gamma)$ and fix $f : X \to X$ an optimal map representing ϕ. Suppose X has an open neighbourhood U such that for any $Y \in U$ obtained from X by a sequence of simple folds directed by f, we have $\lambda(X) \leq \lambda(Y)$. Then there is $Y_n \in U$, all contained in the same simplex, with $Y_n \to X$ and $Y_n \in \mathrm{TT}_0$.*

Proof. The proof is basically the same as in [5]. We recall that for Y obtained from X by folds directed by f, we still denote f the induced optimal map. First we remark that if Y is obtained from X by folds directed by f then $\lambda(Y) \leq \lambda(X)$ and by minimality of X we have $\lambda(Y) = \lambda(X)$. We consider the gate structure induced by f. We call a vertex of $Y_{\max}$ *foldable* if it has at least two elements of $Y_{\max}$ in the same gate.

Locally, by using as small as we want folds in $X_{\max}$, directed by f, we find $Y \in U$ such that

(1) $\lambda(Y) = \lambda(X)$;
(2) Y maximizes the dimension of Δ_Y among points reachable from X via folds directed by f;
(3) Y minimizes $Y_{\max}$ among points of Δ_Y satisfying (1) and (2);
(4) Y maximizes the number of foldable vertices of $Y_{\max}$ among points satisfying (1), (2), (3).

Let $A \subseteq Y_{\max}$ be an f-invariant subgraph given by Lemma 4.18. We claim that $f|_A$ is a strict train track map. Indeed, otherwise there is either an edge e or a legal turn τ in A having illegal image. Let v be the vertex of τ.

- If $f(e)$ contains an illegal turn η then by folding it a little, we would reduce the tension graph, contradicting (3). (Note that $\eta \subset Y_{\max}$ because $A \subseteq Y_{\max}$ is f-invariant, thus by folding η we do not change simplex of $\mathcal{O}(\Gamma)$ because of (2).)
- If $f(\tau)$ is an illegal turn η then we fold it a little. Either $Y_{\max}$ becomes one-gated at v, and in this case the optimization process reduces the tension graph, contradicting (3), or v were not foldable at Y and became foldable, thus contradicting (4).

Finally, note that given such an Y, the sequence Y_n can be chosen in Δ_Y. □

In particular, since $\mathrm{Min}(\phi)$ is clearly closed, we now have:

$$\mathrm{LocMin}(\phi) \subseteq \overline{\mathrm{TT}_0(\phi)} \subseteq \overline{\mathrm{TT}(\phi)} \subseteq \overline{\mathrm{Min}(\phi)} = \mathrm{Min}(\phi) \subseteq \mathrm{LocMin}(\phi)$$

hence all inclusions are equalities.

Lemma 4.20. $\overline{\mathrm{TT}(\phi)} = \mathrm{TT}(\phi)$.

Proof. Let $X \in \overline{\mathrm{TT}(\phi)} = \mathrm{Min}(\phi)$. Let $f : \widetilde{X} \to \widetilde{X}$ be an optimal map representing ϕ. By Lemma 4.19 there is $X_n \to X$ and $f_n \to f$ so that $(X_n, f_n) \in \mathrm{TT}_0(\phi)$. By Lemma 4.14 there is an immersed periodic line

L_n in $(\widetilde{X}_n)_{\max}$ such that $f_n^k(L_n) \subset (\widetilde{X}_n)_{\max}$ is embedded for all $k \in N$. Since the points X_n belong to the same simplex, we can suppose that all the L_n are in fact the same line L. Since $f_n \to f$ and the maps are all PL, $L \subset \widetilde{X}_{\max}$ and $f^k(L) \subset \widetilde{X}_{\max}$. Moreover, if f^k where not injective on L for some k, then we could find $\varepsilon > 0$ and point p, q with $d_X(p,q) = \varepsilon$ and $f^k(p) = f^k(q)$. Now the fact that $f_n \to f$ would contradict the fact that $f_n^k|_L$ is a homothety of ratio $\lambda(\phi)$. Thus $f^k|_L$ is embedded for any k, f is a train track map and so $X \in \mathrm{TT}(\phi)$. □

This completes the proof of Theorem 4.17. □

Corollary 4.21. *Let $\phi \in \mathrm{Out}(\Gamma)$. If $\mathrm{LocMin}(\phi) \neq \emptyset$ then ϕ is either irreducible or thick reducible.*

We end this section by proving a lemma which is basically a rephrasing of Lemma 4.19 with a language which will be more useful in the final part of the paper.

Definition 4.22. Let $\phi \in \mathrm{Out}(\Gamma)$ a point $X \in \mathcal{O}(\Gamma)$ is called an *exit point* of Δ_X if for any neighbourhood U of X in $\mathcal{O}(\Gamma)$ there is a point $X_E \in U$ finite sequence of points $X = X_0, X_1, \ldots, X_m = X_E$ in U, each one obtained by a simple fold directed by an optimal map representing ϕ such that Δ_{X_i} is face at the finite of $\Delta_{X_{i+1}}$, such that Δ_X is a proper face of Δ_{X_E}, and such that

$$\lambda_\phi(X_E) < \lambda_\phi(X)$$

(strict inequality).

Lemma 4.23. *Let $\phi \in \mathrm{Out}(\Gamma)$ and $X \in \mathcal{O}(\Gamma)$ such that $\lambda_\phi(X)$ is a local minimum for λ_ϕ in Δ_X. Suppose $X \notin \mathrm{TT}(\phi)$.*

Then, for any open neighbourhood U of X in Δ_X there is $Z \in U$, obtained from X by folds directed by optimal maps, such that $\lambda_\phi(Z) = \lambda_\phi(X)$, and which admits a simple fold directed by an optimal map and in the tension graph, entering in a simplex Δ' having Δ_X as a proper face. (See Figure 3.)

Moreover, by finitely many such folds we find an X' s.t. Δ_X is a proper face of $\Delta_{X'}$ and $\lambda_\phi(X') < \lambda_\phi(X)$. In particular X is an exit point of Δ_X.

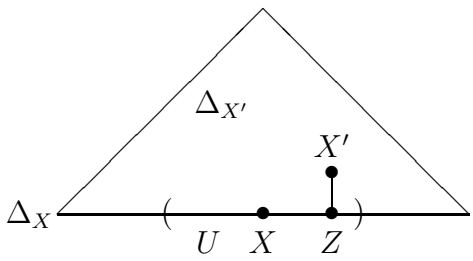

FIGURE 3. Graphical statement of Lemma 4.23

Proof. Let's prove the first claim. Since $X \notin \mathrm{TT}(\phi)$, by Theorem 4.17 there is a neighbourhood of X in Δ_X which is contained in the complement of $\mathrm{TT}_0(\phi)$. Without loss generality we may assume that U is contained in such neighbourhood.

Let $f : X \to X$ be an optimal map representing ϕ. If there is a non-trivalent foldable vertex in $X_{\max}$ then we set $Z = X$ and we are done. Otherwise, consider $Z \in U$ obtained from X by a fold directed by f (we still denote by $f : Z \to X$ the map induced by f). We have $\lambda(Z) \leq \lambda(X)$. Since $\lambda(X)$ is a local minimum, we must have $\lambda(Z) = \lambda(X)$. Let $Y \subset Z_{\max}$ be an f-invariant sub-graph given by Lemma 4.18. Since $Z \notin \mathrm{TT}_0(\phi)$, the restriction $f|_Y$ is not a strict train-track. It follows that by using folds directed by optimal maps we can either

a) reduce the tension graph; or
b) increase the number of foldable vertices; or
c) create a non-trivalent foldable vertex.

So far Z is generic. We choose $Z \in U$ so that, in order:

(1) it locally minimizes the tension graph;
(2) it locally maximizes the number of foldable vertices among points satisfying (1).

For such a Z the only possibility that remains in the above list of alternatives is *c*), and we are done.

The last claim follows from the fact that the dimension of $\mathcal{O}(\Gamma)$ is bounded. □

5. Behaviour of λ at bordification points

For the rest of the section we fix G and $\Gamma = \sqcup_i \Gamma_i$ as in Notation 2.10. We also fix $\phi \in \mathrm{Aut}(\Gamma)$ and we understand that $\lambda = \lambda_\phi$. In this section we discuss the behaviour of λ when we reach points in $\partial_\infty(\mathcal{O}(\Gamma))$. We will see that the function λ is not continuous and we will provide conditions that assure continuity along particular sequences. We will also focus on the behaviour of λ on horospheres. In this section we will mainly think to points of $\mathcal{O}(\Gamma)$ as graphs.

Points near the boundary at infinity have some sub-graph that is almost collapsed. This is usually referred to as the "thin" part of outer space. We will introduce a more fine notion of "thinness".

Definition 5.1. Let $\varepsilon > 0$. A point $X \in \mathcal{O}(\Gamma)$ is *ε-thin* if there is a loop γ in X such that $L_X(\gamma) < \varepsilon \operatorname{vol}(X)$.

Definition 5.2. Let $M, \varepsilon > 0$. A point $X \in \mathcal{O}(\Gamma)$ is *(M, ε)-collapsed* if there is a loop γ in X such that $L_X(\gamma) < \varepsilon \operatorname{vol}(X)$ and for any other loop η such that $L_X(\eta) \geq \varepsilon \operatorname{vol}(X)$ we have $L_X(\eta) > M \operatorname{vol}(X)$.

Definition 5.3. Let $\varepsilon > 0$. For any $X \in \mathcal{O}(\Gamma)$ we define X_ε the ε-thin part of X as the core graph formed by the axes of elements γ with $L_X(\gamma) < \varepsilon \operatorname{vol}(X)$.

Definition 5.4. Let $X \in \mathcal{O}(\Gamma)$. A sub-graph $A \subset X$ is called *ϕ-invariant* if there is a PL-map $f : X \to X$ representing ϕ such that $f(A) \subseteq A$.

We now state some easy facts, the first of which can be found in [3].

Proposition 5.5. *For any $C > \lambda(\phi)$ there is $\varepsilon > 0$ such that for any $X \in \mathcal{O}(\Gamma)$, if $\lambda(X) < C$ and $X_\varepsilon \neq \emptyset$ then X contains a ϕ-invariant subgraph.*

For a proof in the case Γ is connected see [5, Section 8] (connectedness plays in fact no role).

However, we will need a slightly more precise statement, in order to be able to determine a particular invariant subgraph.

Proposition 5.6. *Let $C \geq 1$ and $M > 0$.*
Let $\varepsilon = 1/2 \min\{M/CD, 1/D\}$, where D is the maximal number of (orbits of) edges for any graph in $\mathcal{O}(\Gamma)$. Then, for $X \in \mathcal{O}(\Gamma)$, if $\lambda_\phi(X) < C$ and X is (M, ε)-collapsed, then X_ε is not the whole X and it is ϕ-invariant.

Proof. By definition any edge in X_ε is shorter than $\varepsilon \operatorname{vol}(X)$. Thus we have $\operatorname{vol}(X_\varepsilon) < \varepsilon \operatorname{vol}(X) D$. In particular, since $\varepsilon D < 1$ then $X_\varepsilon \neq X$ (and thus there exists a loop η with $L_X(\eta) > \varepsilon \operatorname{vol}(X)$, whence $L_X(\eta) > M \operatorname{vol}(X)$), since X is (M, ε)-collapsed).

Let $f : X \to X$ be an optimal PL-map representing ϕ.

By picking a maximal tree in the quotient, we may find a generating set of the fundamental group of (each component of) X_ε whose elements have length at most $2 \operatorname{vol}(X_\varepsilon)$. For any such generator, γ, we have that $L_X(f(\gamma))/L_X(\gamma) \leq C$ and hence, $L_X(f(\gamma)) \leq C L_X(\gamma) \leq 2C \operatorname{vol}(X_\varepsilon) < 2CD\varepsilon \operatorname{vol}(X) \leq M \operatorname{vol}(X)$. But since X is (M, ε)-collapsed, we get that $L_X(f(\gamma)) < \varepsilon \operatorname{vol}(X)$. Hence $f(\gamma)$ is homotopic to a loop in X_ε.

Varying γ we deduce that X_ε is ϕ-invariant. □

Proposition 5.7. *Let $X \in \mathcal{O}(\Gamma)$ and $\phi \in \operatorname{Aut}(\Gamma)$. Suppose that $A \subset X$ is a ϕ-invariant core graph. Then $\lambda_{\phi|_A}(A) \leq \lambda_\phi(X)$.*

Proof. Let $f : X \to X$ be a PL-map representing ϕ. Since A is ϕ-invariant, $f(A) \subset A$ up to homotopy. By passing to the universal covering we see that $f|_A : A \to X$ retracts to a map $f_A : A \to A$ representing ϕ with $\operatorname{Lip}(f_A) \leq \operatorname{Lip}(f)$, hence $\lambda_{\phi|_A}(A) \leq \operatorname{Lip}(f_A) \leq \operatorname{Lip}(f) = \lambda_\phi(X)$. □

Theorem 5.8 (Lower semicontinuity of λ)**.** *Fix $\phi \in \operatorname{Aut}(\Gamma)$ and $X \in \mathcal{O}(\Gamma)$. Let $(X_i)_{i \in \mathbb{N}} \subset \Delta_X$ be a sequence such that for any i, $\lambda_\phi(X_i) < C$ for some C. Suppose that $X_i \to X_\infty \in \partial_\infty \Delta_X$ which is obtained from*

X by collapsing a sub-graph $A \subset X$. Then ϕ induces an element of $\mathrm{Aut}(X/A)$*, still denoted by ϕ.*

Moreover $\lambda_\phi(X_\infty) \leq \liminf_{i\to\infty} \lambda_\phi(X_i)$, and if strict inequality holds, then there is a sequence of minimal optimal maps $f_i : X_i \to X_i$ representing ϕ such that eventually on i we have $(X_i)_{\max} \subseteq \mathrm{core}(A)$.

Proof. Let M be the "systole" of X_∞, that is to say the shortest length of simple non-trivial loops in X_∞. For any $M/\operatorname{vol}(X) > \varepsilon > 0$, eventually on i, X_i is $(M/2\operatorname{vol}(X), \varepsilon)$-collapsed and $(X_i)_\varepsilon = \mathrm{core}(A)$. By Proposition 5.6 A is ϕ-invariant, thus $\phi \in \mathrm{Aut}(X/A)$.

For any loop γ the lengths $L_{X_i}(\gamma)$ and $L_{X_i}(\phi(\gamma))$ converge to $L_{X_\infty}(\gamma)$ and $L_{X_\infty}(\phi(\gamma))$ respectively. Therefore, if γ is a candidate in X_∞ that realizes $\lambda_\phi(X_\infty)$, we have that $\lambda_\phi(X_i) \geq L_{X_i}(\phi(\gamma))/L_{X_i}(\gamma) \to \lambda_\phi(X_\infty)$ whence the lower semicontinuity of λ.

On the other hand, by Theorem 3.23 for any i there is a minimal optimal map $f_i : X_i \to X_i$ representing ϕ. Let γ_i be a candidate that realizes $\lambda_\phi(X_i)$, i.e. a f_i-legal candidate in $(X_i)_{\max}$. Since X is combinatorically finite, we may assume w.l.o.g. that $\gamma_i = \gamma$ is the same loop for any i. We have

$$\lambda_\phi(X_i) = \frac{L_{X_i}(\phi(\gamma))}{L_{X_i}(\gamma)} \to \frac{L_{X_\infty}(\phi(\gamma))}{L_{X_\infty}(\gamma)}$$

Thus if $L_{X_\infty}(\gamma) \neq 0$ we have $\lambda_\phi(X_\infty) = \liminf \lambda_\phi(X_i)$. It follows that if there is a jump in λ at X_∞, then any legal candidate is contained in A. Since f_i is minimal this implies that $\mathrm{core}(A)$ contains the whole tension graph. □

Remark 5.9. A comment on Theorem 5.8 is required. To avoid cumbersome notation, we have decided to denote by ϕ both the element of $\mathrm{Aut}(X)$ and the one induced in $\mathrm{Aut}(X/A)$. So when we write $\lambda_\phi(X_\infty)$ we mean $\Lambda(X_\infty, \phi X_\infty)$ as elements in $\mathcal{O}(X/A)$. In particular, $\lambda_\phi = \inf_X \lambda_\phi(X)$ can be different if computed in $\mathcal{O}(X)$ or in $\mathcal{O}(X/A)$. When this will be crucial we will specify in which space we take the infimum.

Moreover, if $\phi|_A$ is the restriction of ϕ to A, then $\lambda_{\phi|_A}$ is calculated in the space $\mathcal{O}(A)$. While the simplex Δ_{X_∞} is a simplicial face of Δ_X, $\Delta_A \in \mathcal{O}(A)$ has not that meaning. One could argue that Δ_A is the simplex "opposite" to Δ_{X_∞} in Δ, but ϕ does not necessarily produces an element of $\mathrm{Aut}(X/(X \setminus A))$ as the complement of A may be not invariant.

Clearly, if $A \subset X$ is ϕ-invariant then $\lambda_\phi(X/A) < \infty$. On the other hand, if A is not ϕ-invariant, its collapse makes λ explode. Thus we can extend the function λ as follows.

Definition 5.10. Let $X_\infty \in \partial_\infty \mathcal{O}(\Gamma)$. We say that $\lambda_\phi(X_\infty) = \infty$ if X_∞ is obtained from a Γ-graph X by collapsing a core sub-graph $A \subset X$ which is not ϕ-invariant.

In general, the function λ is not uniformly continuous with respect to the Euclidean metric, even in region where it is bounded, and so we cannot extend it to the simplicial closure of simplices. However we see now that the behaviour of λ is controlled on segments.

We recall the description of horospheres given in 2.5. Suppose that X_∞ is obtained from a Γ-graph X by collapsing a ϕ-invariant core subgraph $A = \cup_i A_i$. Let k_i be the number of germs of edges incidents to A_i in $X \setminus A$. Then $\mathrm{Hor}(X_\infty)$ is a product of outer spaces with marked points $\mathcal{O}(A_i, k_i)$.

Notation 5.11. We denote $\pi : \mathrm{Hor}(X_\infty) \to \mathbb{P}\mathcal{O}(A)$ the projection that forgets marked points.

Note that we chosen X_∞ not projectivized and $\mathbb{P}\mathcal{O}(A)$ projectivized. For any $Y \in \mathbb{P}\mathcal{O}(A)$ if $Z \in \pi^{-1}(Y)$, then there is a scaled copy of Y in Z. We denote by $\mathrm{vol}_A(Y)$ the volume of Y in Z. With this notation in place, we can now prove a key regeneration lemma.

Lemma 5.12 (Regeneration of optimal maps)**.** *Fix $\phi \in \mathrm{Aut}(\Gamma)$ ad $X \in \mathcal{O}(\Gamma)$. Let $X_\infty \in \partial_\infty \Delta_X$ be obtained from X by collapsing a ϕ-invariant core sub-graph A. Then, for any* PL*-map $f_A : A \to A$ representing $\phi|_A$, and for any $\varepsilon > 0$ there is $X_\varepsilon \in \Delta_X$ such that*

$$\lambda_\phi(X_\varepsilon) \leq \max\{\lambda_\phi(X_\infty) + \varepsilon, \mathrm{Lip}(f_A)\}.$$

More precisely, for any $Y \in \mathbb{P}\mathcal{O}(A)$ and map $f_Y : Y \to Y$ representing $\phi|_A$, for any map $f : X_\infty \to X_\infty$ representing ϕ, for any $\widehat{X} \in \mathrm{Hor}(X_\infty) \cap \pi^{-1}(Y)$, and for any $\varepsilon > 0$; there is $0 < \delta = \delta(f, f_Y, X_\infty, \Delta_{\widehat{X}})$, such that for any $Z \in \Delta_{\widehat{X}} \cap \pi^{-1}(Y)$, if $\mathrm{vol}_Z(Y) < \delta$ there is a PL*-map $f_Z : Z \to Z$ representing ϕ such that $f_Z = f_Y$ on Y and*

$$\mathrm{Lip}(f_Z) \leq \max\{\lambda_\phi(X_\infty) + \varepsilon, \mathrm{Lip}(f_Y)\}$$

(and hence the optimal map $\mathrm{opt}(f_z)$ satisfies the same inequality[6]).

Proof. For this proof we will need to work with both graphs and tree, and we will the usual $\tilde{}$-notation for the universal coverings. We denote by $\sigma : X \to X_\infty$ the map that collapses A. If A_i is a component of A, we denote by v_i the non-free vertex $\sigma(A_i)$. We set $V_A = \{v_i\}$. Let k_i be the valence of v_i in X_∞. For any v_i let $E_i^1, \ldots, E_i^{k_i}$ be the half-edges incident to v_i in X_∞. Let Y_i be the components of Y. Points in $\Delta_{\widehat{X}}$ are built by inserting a scaled copy of each Y_i at the v_i as follows. (Now we need to pass to the universal coverings.)

For every half-edge E_i^j of X_∞ we choose a lift in $\widetilde{X}_\infty$. The tree $\widetilde{\widehat{X}}$ is given by attaching $\widetilde{E}_i^j$ to a point $\widetilde{y}_i^j$ of $\widetilde{Y}_i$, and then equivariantly attaching any other lift of the E_i^j. At the level of graphs this is equivalent to choose $y_i^j \in Y_i$. Two different choices at the level of universal coverings

[6]We notice that while $f_Z = f_Y$ on Y, this is no longer true for $\mathrm{opt}(f_Z)$

differ, at the level of graphs by, closed paths γ_i^j in Y_i and based at y_i^j. The choice of the simplex $\Delta_{\widehat{X}}$ fixes such ambiguity. Moreover for any two graphs in $\pi^{-1}(Y) \cap \Delta_{\widehat{X}}$ the points y_i^j are attached to the the same edge of Y_i. Let $Z \in \pi^{-1}(Y) \cap \Delta_{\widehat{X}}$.

Given $f_Y : Y \to Y$, consider its lift to $\widetilde{Y}$ and set $\widetilde{z}_i^j = \widetilde{f}_Y(\widetilde{y}_i^j)$. There is a unique embedded arc $\tilde{\gamma}_i^j$ from $\widetilde{z}_i^j$ to $\widetilde{y}_i^j$. Let L_i be the simplicial length of γ_i^j. L_i depends only on f_Y and the choices of $\widetilde{y}_i^j$, hence it depends only on f_Y and $\Delta_{\widehat{X}}$.

Now, given $f : X_\infty \to X_\infty$, there exists a continuous map $g : Z \to Z$ representing ϕ, which agrees with f_Y on Y and which is obtained by a perturbation of f on edges of X_∞. Namely on E_i^j we need to attach γ_i^j to $f(E_i^j)$, and in each point of $f^{-1}(V_A)$ we need to insert a small segment whose image is a suitable path in Y. We refer the reader to the Appendix (section 12) for an accurate and detailed discussion on the properties of such a map. For the present purpose it is sufficient to note that there is a constant C such that g can be obtained so that $\mathrm{Lip}(\mathrm{PL}(g)) < \mathrm{Lip}(f) + C\,\mathrm{vol}(Y)$. Moreover the constant C depends only on the L_i's, the paths added in $f^{-1}(V_A)$, and the edge-lengths of X_∞. Hence it depends only on $f_Y, \Delta_{\widehat{X}}, X_\infty$.

The thesis follows by setting $\delta < \varepsilon/C$ and $f_Z = \mathrm{PL}(g)$. □

Definition 5.13. Fix $\phi \in \mathrm{Aut}(\Gamma)$. Let $X_\infty \in \partial_\infty \Delta \subset \mathcal{O}(\Gamma)$. We say that X_∞ has *not jumped in* Δ if there is a sequence of points $X_i \in \Delta$ such that $\lambda_\phi(X_\infty) = \lim_i \lambda_\phi(X_i)$. We say that $X_\infty \in \partial_\infty \mathcal{O}(\Gamma)$ *has not jumped* if there is a simplex $\Delta \subset \mathrm{Hor}(\Delta_{X_\infty})$ such that X_∞ has not jumped in Δ.

The above definition is for points in $\partial_\infty \mathcal{O}(\Gamma)$. We decide to say that X has not jumped for any $X \in \mathcal{O}(\Gamma)$.

Notice that even if X_∞ has not jumped, there may exist a simplex $\Delta \in \mathrm{Hor}(\Delta_{X_\infty})$ such that X_∞ has jumped in Δ. This is because if A is the collapsed part and $\phi|_A$ does not have polynomial growth, then we can choose a point in $\mathcal{O}(A)$ with arbitrarily high $\lambda_{\phi|_A}$. Moreover, even if X_∞ has not jumped in Δ it may happen that X_∞ is not a continuity point of λ. For example if the collapsed part A has a sub-graph B which is not invariant, then the collapse of B forces λ to increase due to Proposition 5.6, and thus we can approach X_∞ with arbitrarily high λ.

Theorem 5.14. *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let $X \in \mathcal{O}(\Gamma)$ containing an invariant sub-graph A. Let $X_\infty = X/A$ and $C = \mathrm{core}(A)$. Then*

$$\lambda_{\phi|_C}(\Delta_C) \leq \lambda_\phi(\Delta_X).$$

Moreover X_∞ has not jumped in Δ_X if and only if

$$\lambda_\phi(X_\infty) \geq \lambda_{\phi|C}(\Delta_C)$$

if and only if

$$\lambda_\phi(X_\infty) \geq \lambda_\phi(\Delta_X).$$

In particular there is gap as $\lambda_\phi(X_\infty)$ cannot belong to the interval $(\lambda_{\phi|_C}(\Delta_C), \lambda_\phi(\Delta_X))$ (if non-empty).

Proof. The first claim is a direct consequence of Proposition 5.7. Let $X_i \in \Delta_X$ with $X_i \to X_\infty$ without jump, and let A_i be the metric version of A in X_i. Then by Proposition 5.7 we have

$$\lambda_{\phi|_C}(\Delta_C) \leq \lambda_{\phi|_C}(\mathrm{core}(A_i)) \leq \lambda_\phi(X_i) \to \lambda_\phi(X_\infty).$$

Conversely, suppose $\lambda_\phi(X_\infty) \geq \lambda_{\phi|C}(\Delta_C)$. For any $\varepsilon > 0$ there is $C_\varepsilon \in \Delta_C$ and a PL-map $f_{C_\varepsilon} : C_\varepsilon \to C_\varepsilon$ representing $\phi|_C$ such that $\mathrm{Lip}(f_{C_\varepsilon}) < \lambda_{\phi|_C}(\Delta_C) + \varepsilon$. By Lemma 5.12 there is a point $X_\varepsilon \in X$ and a map $f_\varepsilon : X_\varepsilon \to X_\varepsilon$ representing ϕ such that $X_\varepsilon \to X_\infty$ as $\varepsilon \to 0$ and $\mathrm{Lip}(f_e) \leq \lambda_\phi(X_\infty)$. Then $\lambda_\phi(X_\varepsilon) \to \lambda_\phi(X_\infty)$. The second claim is proved. Now, if $\lambda_\phi(X_\infty) \geq \lambda(\Delta_X)$, by the first two claims it has not jumped in Δ_X. And if X_∞ as not jumped

$$\lambda(\Delta_X) \leq \lambda_\phi(X_i) \to \lambda_\phi(X_\infty).$$

□

Lemma 5.15. *Let $\phi \in \mathrm{Aut}(\Gamma)$ and let $X \in \mathcal{O}(\Gamma)$. If $\lambda(\phi) > 1$, then λ_ϕ is not bounded on Δ_X.*

Proof. If there is a loop which is not ϕ-invariant, then by collapsing that loop we force λ_ϕ to explode. If any loop is ϕ-invariant then by Theorem 3.7 we get $\lambda(\phi) = 1$. □

Lemma 5.16. *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let $X \in \mathcal{O}(\Gamma)$ containing an invariant sub-graph A. Let $X_\infty = X/A$ and let $C = \mathrm{core}(A)$. Let*

$$X_t = (1-t)X_\infty + tX$$

and let C_t be the metric version of C in X_t. If $\lambda_\phi(X_\infty) < \liminf \lambda_\phi(X_t)$ then then for $t > 0$ small enough $\lambda_\phi(X_t)$ is locally constant, more precisely we have

$$\lambda_\phi(X_t) = \lambda_{\phi|_C}(\mathrm{core}(A_1)).$$

In particular, this is the case if X_∞ has jumped in Δ along the segment XX_∞.

Proof. By Lemma 5.8 for t small enough there is an optimal map $f_t : X_t \to X_t$ whose tension graph of X_t is contained in C_t. Since C_t is ϕ-invariant, $f_t(C_t) \subset C_t$ up to homotopy. Since the vertices of $(X_t)_{\max}$ are at least two gated, $f((X_t)_{\max}) \subset C_t$. Therefore $\lambda_{\phi|_C}(C_t) = \mathrm{Lip}(f_t)$ and $\lambda_\phi(X_t) = \mathrm{Lip}(f_t) = \lambda_{\phi|_C}(C_t) = \lambda_{\phi|_C}(C_1)$ (where the last equality follows from the fact that $[C_t] = [C_1] \in \mathbb{P}\mathcal{O}(C)$).

The last claim follows because by Theorem 5.14, and since X_∞ has jumped in Δ, we have

$$\lambda_\phi(X_\infty) < \lambda_{\phi|_C}(\Delta_C) \leq \lambda_\phi(\Delta) \leq \lambda_\phi(X_t)$$

hence $\lambda_\phi(X_\infty) < \liminf_t \lambda_\phi(X_t)$. □

Corollary 5.17. *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let Δ be a simplex of $\mathcal{O}(\Gamma)$. If there a point in $\overline{\Delta}^\infty$ which jumped in Δ, then there is a min-point in Δ that realizes $\lambda(\Delta)$.*

In particular, there is always a min-point $X_{\min}$ in $\overline{\Delta}^\infty$ which has no jumped and such that

$$\lambda_\phi(\Delta_X) = \lambda_\phi(X_{\min}) = \lambda_\phi(\Delta_{X_{\min}}).$$

Proof. Let $\Delta = \Delta_X$ for some Γ-graph X. Suppose that A is a ϕ-invariant subgraph so that $X_\infty = X/A$ has a jump in Δ. Let $C = \mathrm{core}(A)$. By Theorem 5.14

$$1 \leq \lambda(X_\infty) < \lambda(\Delta_C) \leq \lambda(\Delta)$$

hence by Lemma 5.15 there $C_o \in \Delta_C$ be such that $\lambda(C_o) = \lambda(\Delta)$. Thus there is $X_o \in \Delta$ obtained by inserting on X_∞ a copy of a metric graph isomorphic to A and with core C_0. By Lemma 5.16, for small enough t, we have $\lambda(tX + (1-t)X_\infty) = \lambda(\Delta)$. The last claim follows clearly from the first. □

6. Convexity properties of the displacement function

We recall that we are using the terminology "simplex" in a wide sense, as Δ_X is a standard simplex if we work in $\mathbb{P}\mathcal{O}(\Gamma)$ and the cone over it if we work in $\mathcal{O}(\Gamma)$. (Remember we use Notation 2.10 for Γ.)

The function λ is scale invariant on $\mathcal{O}(\Gamma)$ so it descends to a function on $\mathbb{P}\mathcal{O}(\Gamma)$. In order to control the value of λ on segments in terms of its value on vertices, we would like to say that λ is convex on segments. A little issue appears with projectivization. If Δ is a simplex of $\mathcal{O}(\Gamma)$, its euclidean segments are well defined, and their projections on $\mathbb{P}(\mathcal{O}(\Gamma))$ are euclidean segments in the image of Δ. However, the linear parametrization is not a projective invariant (given X, Y, the points $(X+Y)/2$ and $(5X+Y)/2$ are in different projective classes).

It follows that convexity of a scale invariant function is not well-defined. In fact if σ is a segment in Δ, $\pi : \Delta \to \mathbb{P}\Delta$ is the projection, and f is a convex function on σ, then $f \circ \pi^{-1}$ may be not convex. It is convex only up to reparametrization of the segment $\pi(\sigma)$. Such functions are called quasi-convex, and this notion will be enough for our purposes.

Definition 6.1. A function $f : [A, B] \to \mathbb{R}$ is called *quasi-convex* if for all $[a, b] \subseteq [A, B]$

$$\forall t \in [a, b] \qquad f(t) \leq \max\{f(a), f(b)\}.$$

Note that quasi-convexity is scale invariant.

Lemma 6.2. *For any $\phi \in \mathrm{Aut}(\Gamma)$ and for any open simplex Δ in $\mathcal{O}(\Gamma)$ the function λ is quasi-convex on segments of Δ. Moreover, if $\lambda(A) > \lambda(B)$ then λ is not locally constant near A.*

Proof. Let X be a Γ-graph such that $\Delta = \Delta_X$. We use the Euclidean coordinates of Δ labelled with edges of X, namely a point P in Δ is given by a vector whose e^{th} entry is the length of the e in P. In the same way, to any reduced loop η in X we associate its occurrence vector, whose e^{th} entry is the number of times that η passes through the edge e. We will denote by η both the loop and its occurrence vector. With this notation, the length function is bilinear:

$$L_X(\gamma) = \langle X, \gamma \rangle$$

(where $\langle , \rangle$ denotes the standard scalar product on $\mathbb{R}^k$.)

Let σ be a segment in Δ with endpoints A, B. Let γ be a candidate. We consider both γ and $\phi\gamma$ as loops in X. Up to switching A and B, we may assume that

$$\frac{\langle A, \phi\gamma \rangle}{\langle A, \gamma \rangle} \geq \frac{\langle B, \phi\gamma \rangle}{\langle B, \gamma \rangle}.$$

Such a condition is scale invariant, and since λ is scale invariant, up to rescaling B we may assume that $\langle B, \gamma \rangle > \langle A, \gamma \rangle$. Now, we parametrize σ as usual with $[0, 1]$

$$\sigma(t) = A_t = Bt + (1-t)A$$

We are interested in the function

$$F_\gamma(t) = \frac{\langle A_t, \phi\gamma \rangle}{\langle A_t, \gamma \rangle} = \frac{\langle Bt + (1-t)A, \phi\gamma \rangle}{\langle Bt + (1-t)A, \gamma \rangle} = \frac{\langle A, \phi\gamma \rangle + t\langle B - A, \phi\gamma \rangle}{\langle A, \gamma \rangle + t\langle B - A, \gamma \rangle}$$

A direct calculation shows that the second derivative of a function of the type $f(t) = (a + tb)/(c + td)$ is given by $2(ad - bc)d/(c + td)^3$.

So the sign of $F''_\gamma(t)$ is given by

$$\big(\langle A, \phi\gamma \rangle \langle B, \gamma \rangle - \langle B, \phi\gamma \rangle \langle A, \gamma \rangle\big)\big(\langle B - A, \gamma \rangle\big)$$

which is positive by our assumption on A, B. Hence $F_\gamma(t)$ is convex and therefore quasi-convex:

$$F_\gamma(t) \leq \max\{F_\gamma(A), F_\gamma(B)\}.$$

Now, by the Sausage Lemma 3.7 we have:

$$\begin{aligned} \lambda_\phi(A_t) &= \max_\gamma F_\gamma(t) \leq \max\{\max_\gamma F_\gamma(A), \max_\gamma F_\gamma(B)\} \\ &= \max\{\lambda_\phi(A), \lambda_\phi(B)\}. \end{aligned}$$

Finally, since lengths of candidates are finitely many, there is a candidate γ_o such that for t sufficiently small we have $\lambda_\phi(A_t) = F_{\gamma_o}(t)$. So, if λ_ϕ is locally constant near A, then we must have $F''_{\gamma_o}(t) = 0$ hence

$$\lambda_\phi(A) = \lambda_\phi(A_0) = \frac{\langle A, \phi\gamma_o \rangle}{\langle A, \gamma_o \rangle} = \frac{\langle B, \phi\gamma_o \rangle}{\langle B, \gamma_o \rangle} \leq \lambda_\phi(B).$$

□

Lemma 6.3. *Let $\phi \in \mathrm{Aut}(\Gamma)$ and let Δ be a simplex in $\mathcal{O}(\Gamma)$. Let $A, B \in \overline{\Delta}^{\infty}$ be two points that have not jumped in Δ. Then for any $P \in \overline{AB}$*

$$\lambda(P) \leq \max\{A, B\}$$

Moreover, if $\lambda(A) \geq \lambda(B)$, then $\lambda|_{\overline{AB}}$ is continuous at A.

Proof. Let X be a graph of groups so that $\overline{AB} = \Delta_X$. By Lemma 6.2, the function λ is quasi-convex on the interior of $\overline{AB}$ as a segment in $\mathcal{O}(X)$. Let $\{A_i\}$ and $\{B_i\}$ sequences in Δ such that $A = \lim A_i$ and $B = \lim B_i$ with $\lim \lambda(A_i) = \lambda(A)$ and $\lim \lambda(B_i) = \lambda(B)$. Such sequences exists because of the non jumping hypothesis. For all point P in the segment $\overline{AB}$, there is a sequence of points P_i in the segment $\overline{A_iB_i}$ such that $P_i \to P$. By Lemma 6.2 we know

$$\lambda(P_i) \leq \max\{\lambda(A_i), \lambda(B_i)\},$$

and by lower semicontinuity (Theorem 5.8) of λ and the non jumping assumption, such inequality passes to the limit. In particular, if $\lambda(A) \geq \lambda(B)$, then $\lambda(P) \leq \lambda(A)$ for any $P \in \overline{AB}$.

Now suppose that $P^j \to A$ is a sequence in the segment $\overline{AB}$. Let $P^j = \lim_i P_i^j$. Then by lower semicontinuity Theorem 5.8 applied to the space $\mathcal{O}(P)$, on the segment $\overline{AB}$ we have

$$\lambda(A) \geq \lim_j \lambda(P_j) \geq \lambda(A).$$

□

We end this section with an estimate of the derivative of functions like the $F_\gamma(t)$ defined as in Lemma 6.2, which will be used in the sequel. As above, we use the formalism $\langle X, \gamma\rangle = L_X(\gamma)$.

Lemma 6.4. *Let $\Delta = \Delta_X$ be a simplex of $\mathcal{O}(\Gamma)$ and $A, B \in \overline{\Delta}^{\infty}$. Let γ be a loop in X which is not collapsed neither in A nor in B and set*

$$C = \max\{\frac{L_A(\gamma)}{L_B(\gamma)}, \frac{L_B(\gamma)}{L_A(\gamma)}\}$$

Let ϕ be any automorphism of Γ. Suppose that $\frac{\langle B,\phi\gamma\rangle}{\langle B,\gamma\rangle} \geq \frac{\langle A,\phi\gamma\rangle}{\langle A,\gamma\rangle}$. Let $A_t = tB + (1-t)A$ be the linear parametrization of the segment AB in Δ and define $F_\gamma(t) = \frac{\langle A_t,\phi\gamma\rangle}{\langle A_t,\gamma\rangle}$. Then

$$0 \leq F'_\gamma(t) \leq C\frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle}$$

In particular, for any point P in the segment AB we have

$$\lambda_\phi(P) \geq \frac{\langle P, \phi\gamma\rangle}{\langle P, \gamma\rangle} \geq \frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle} - C\lambda_\phi(B)\frac{||P - B||}{||A - B||}$$

where $||X - Y||$ denotes the standard Euclidean metric on Δ.

Before the proof, a brief comment on the statement is desirable. First, note that the constant C does not depend on ϕ. Moreover, by taking the supremum where γ runs over all candidates given by the Sausage Lemma 3.7, then C does not even depend on γ. Finally if γ is a candidate that realizes $\lambda_\phi(B)$, then we get a bound of the steepness of F_γ which does not depend on ϕ nor on γ but just on $\lambda_\phi(B)$ and $||A-B||$.

Proof. We have

$$F_\gamma(t) = \frac{\langle A_t, \phi\gamma\rangle}{\langle A_t, \gamma\rangle} = \frac{\langle Bt + (1-t)A, \phi\gamma\rangle}{\langle Bt + (1-t)A, \gamma\rangle} = \frac{\langle A, \phi\gamma\rangle + t\langle B-A, \phi\gamma\rangle}{\langle A, \gamma\rangle + t\langle B-A, \gamma\rangle}$$

and a direct calculation show that

$$F'_\gamma(t) = \frac{\langle B, \gamma\rangle\langle A, \gamma\rangle}{(\langle A_t, \gamma\rangle)^2}\left(\frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle} - \frac{\langle A, \phi\gamma\rangle}{\langle A, \gamma\rangle}\right) \tag{1}$$

The first consequence of this equation is that the sign of F'_γ does not depend on t, and since $\frac{\langle B,\phi\gamma\rangle}{\langle B,\gamma\rangle} \geq \frac{\langle A,\phi\gamma\rangle}{\langle A,\gamma\rangle}$, then $F'_\gamma \geq 0$. Moreover, since $\langle A_t, \gamma\rangle$ is linear on t, we have $\frac{\langle B,\gamma\rangle\langle A,\gamma\rangle}{(\langle A_t,\gamma\rangle)^2} \leq C$. Therefore we get

$$F'_\gamma(t) \leq C\frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle}$$

and the first claim is proved. For the second claim, note that the parameter t is nothing but $||A - A_t||/||A - B||$ and thus

$$F_\gamma(1) - F_\gamma(t) \leq (1-t)C\frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle} = \frac{||B - A_t||}{||B - A||}C\frac{\langle B, \phi\gamma\rangle}{\langle B, \gamma\rangle}.$$

If $P = A_t$, we have $F_\gamma(1) = \frac{\langle B,\phi\gamma\rangle}{\langle B,\gamma\rangle}$ and $F_\gamma(t) = \frac{\langle P,\phi\gamma\rangle}{\langle P,\gamma\rangle}$. By taking in account $\lambda_\phi(B) \geq \frac{\langle B,\phi\gamma\rangle}{\langle B,\gamma\rangle}$ and $\lambda_\phi(P) \geq \frac{\langle P,\phi\gamma\rangle}{\langle P,\gamma\rangle}$ we get the result. □

7. Existence of minimal displaced points and train tracks at the bordification

The existence of points that minimizes the displacement is proved in [5, Theorem 8.4] for irreducible automorphisms, but in fact the philosophy of the proof works in the general case if we are allowed to pass to the boundary at infinity, and by taking in account possible jumps. A part jumps, the problem is that one cannot use compactness for claiming that a minimizing sequence has a limit as the bordification of $\mathcal{O}(\Gamma)$ is not even locally compact. The trick used in [5] is to use Sausage Lemma. We use Notation 2.10 for G and Γ.

Lemma 7.1. *For any Γ, for any $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ the set $\{\lambda_\phi(X) : \phi \in \mathrm{Out}(\Gamma)\}$ is discrete. In other words, given X, all possible displacements of X with respect to all automorphisms (and hence markings) run over a discrete set.*

Proof. This proof is similar to that of [5, Theorem 8.4], we include it by completeness. If $\phi \notin \mathrm{Aut}(X)$ (i.e. if X has a collapsed part which is not ϕ-invariant) than $\lambda_\phi(X) = \infty$ and there is nothing to prove. Otherwise, by Sausage Lemma 3.7, $\lambda_\phi(X) = \Lambda(X, \phi X)$ is computed by the quotient of translation lengths of candidates. The possible values of $L_X(\phi\gamma)$ (with γ any loop) are a discrete set just because X has finitely many edges. Candidates are in general infinitely many, but there are only finitely many lengths of them. Thus the possible values of $\Lambda(X, \phi X)$ runs over a discrete subset of $\mathbb{R}$. □

Theorem 7.2. *For any Γ the global simplex-displacement spectrum*

$$\mathrm{spec}(\Gamma) = \left\{\lambda_\phi(\Delta) : \Delta \text{ a simplex of } \overline{\mathcal{O}(\Gamma)}^\infty, \phi \in \mathrm{Out}(\Gamma)\right\}$$

is well-ordered as a subset of $\mathbb{R}$. In particular, for any $\phi \in \mathrm{Out}(\Gamma)$ the spectrum of possible minimal displacements

$$\mathrm{spec}(\phi) = \left\{\lambda_\phi(\Delta) : \Delta \text{ a simplex of } \overline{\mathcal{O}(\Gamma)}^\infty\right\}$$

is well-ordered as a subset of $\mathbb{R}$.

Proof. Recall that we defined $\lambda_\phi(\Delta)$ as $\inf_{X\in\Delta} \lambda_\phi(X)$. For this proof we fix the volume-one normalization, and in any simplex we use the standard Euclidean norm, denoted by $||\cdot||$.

We argue by induction on the rank of Γ (See Definition 2.11). Clearly if the rank of Γ is one there is nothing to prove. We now assume the claim true for any Γ' of rank smaller than Γ.

We will show than any monotonically decreasing sequence in $\mathrm{spec}(\Gamma)$ has a (non trivial) sub-sequence which is constant, whence the original sequence is eventually constant itself. This implies that $\mathrm{spec}(\Gamma)$ is well-ordered. For the second claim, since $\mathrm{spec}(\phi)$ is a subset of a well-ordered set, it is well-ordered.

We follow the line of reasoning of [5, Theorem 8.4]. Let $\lambda_i \in \mathrm{spec}(\Gamma)$ be a monotonically decreasing sequence. Note that displacements are non-negative so λ_i converges. For any i we chose ϕ_i and a point $X_i \in \overline{\mathcal{O}(\Gamma)}^\infty$ such that $\lambda_{\phi_i}(X_i) = \lambda_{\phi_i}(\Delta_{X_i}) = \lambda_i$ and let $\Delta_i = \Delta_{X_i}$. Up to possibly passing to sub-sequences we may assume that there is $\psi_i \in \mathrm{Out}(\Gamma)$ such that $\psi_i X_i$ belongs to a fixed simplex Δ. Therefore, by replacing ϕ_i with $\psi_i\phi_i\psi_i^{-1}$ we may assume that the X_i all belong to the same simplex Δ. Let X be the graph of groups corresponding to Δ, i.e. $\Delta = \Delta_X$.[7]

Up to sub-sequences, X_i converges to a point X_∞ in the simplicial closure of Δ. By Lemma 7.1 up to possibly passing to a subsequence we may assume that $\lambda_{\phi_i}(X_\infty)$ is a constant L. (The only issue here is that X_∞ is displaced a finite amount by each ϕ_i - we show this is true for all but finitely many i.)

[7]Note that X may be a boundary point of $\mathcal{O}(\Gamma)$ and that we have made no assumption about jumps, so X may jump.

Note that if X_∞ is in $\partial_\infty \overline{\Delta}$, then there exist $M, \varepsilon > 0$ such that X_i is eventually (M, ε)-collapsed. Namely, assuming X_∞ has volume 1, take M to be the length of the shortest loop in X_∞, and take ϵ to be a constant small enough to satisfy the hypotheses of Proposition 5.6. Then the thin part, $(X_i)_\varepsilon$, is the core of the the sub-graph of X which collapsed to obtain X_∞. By Proposition 5.6 $(X_i)_\varepsilon$ is ϕ_i-invariant and so $\lambda_{\phi_i}(X_\infty) < \infty$. Hence $L < \infty$.

Since X_i is a min-point for the function λ_{ϕ_i}, by Lemma 6.2 the function λ_{ϕ_i} either is constant on the segment $X_i X_\infty$ or it is not locally constant near X_∞. By Lemma 5.16 in the latter case X_∞ has not jumped w.r.t. λ_{ϕ_i} along the segment $X_i X_\infty$.

Therefore we have the following three cases, and up to subsequences we may assume that we are in the same case for any i:

(1) λ_{ϕ_i} is constant and continuous on $X_i X_\infty$;
(2) λ_{ϕ_i} is constant on the interior of $X_i X_\infty$ and there is a jump at X_∞, hence $\lambda_{\phi_i}(X_\infty) < \lambda_{\phi_i}(X_i)$ by lower semicontinuity Theorem 5.8;
(3) λ_{ϕ_i} is monotone increasing near X_∞ and continuous at X_∞.

In the first case $\lambda_i = \lambda_{\phi_i}(X_i) = L$ and we are done. In the second case we use the inductive hypothesis. Namely, by Lemma 5.16 there is a core ϕ_i-invariant sub-graph A_i of X_i such that $\lambda_{\phi_i}(X) = \lambda_{\phi_i|_{A_i}}(A_i)$ for any X in the interior of the segment $X_i X_\infty$. Moreover, up to sub-sequences we may assume that A_i is topologically the same graph for any i. If A_i does not minimizes locally $\lambda_{\phi_i|_{A_i}}$ in its simplex, the we could perturb a little A_i and strictly decrease $\lambda_{\phi_i}(X_i)$ contradicting the minimality of X_i. By quasi-convexity Lemma 6.2, in any simplex local minima are global minima and thus $\lambda_{\phi_i|_{A_i}}(A_i) = \lambda_{\phi_i|_{A_i}}(\Delta_{A_i})$. By induction the global simplex-displacement spectrum of A_i is well ordered, hence the decreasing sequence $\lambda_i = \lambda_{\phi_i}(X_i) = \lambda_{\phi_i|_{A_i}}(\Delta_{A_i})$ is eventually constant.

It remains case (3). In this case

$$\lambda_{\phi_i}(X_i) < L = \lambda_{\phi_i}(X_\infty).$$

Let $R > 0$ be such that for any face Δ' of Δ such that $X_\infty \notin \overline{\Delta'}^\infty$, the ball $B(X_\infty, 2R)$ is disjoint from Δ'. In other words, if $P \in B(X_\infty, 2R)$ is obtained form X by collapsing a sub-graph P_0, then P_0 is collapsed also in X_∞. Eventually on i, $X_i \in B(X_\infty, R)$. Let Y_i be the point on the Euclidean line trough X_i, X_∞ at distance exactly R from X_∞.

Let γ_i be a candidate in X_∞ that realizes $\lambda_{\phi_i}(X_\infty)$ and such that the stretching factor $\frac{L_X(\phi_i(\gamma_i))}{L_X(\gamma_i)}$ of γ_i locally decreases toward X_i. Such a γ_i exists because $\lambda_{\phi_i}(X_i) < \lambda_{\phi_i}(X_\infty)$.

By Lemma 6.4 applied with $A = Y_i$ and $B = X_\infty$ we have

$$\lambda_{\phi_i}(X_i) \geq \lambda_{\phi_i}(X_\infty)\left(1 - C\frac{||X_i - X_\infty||}{R}\right).$$

where $C = \max\{\frac{L_{Y_i}(\gamma_i)}{L_{X_\infty}(\gamma_i)}, \frac{L_{X_\infty}(\gamma_i)}{L_{Y_i}(\gamma_i)}\}$. Since there are finitely many lengths of candidates and by our choice of R, the constant C is uniformly bounded independently on i. Since $X_i \to X_\infty$ we have $\varepsilon_i = ||X_i - X_\infty|| \to 0$ and thus

$$L(1 - C\varepsilon_i) \leq \lambda_{\phi_i}(X_i) \leq L.$$

Thus $\lambda_i \to L$ and since it is a monotonically decreasing sequence bounded above by its limit, it must be constant. □

We suspect that $\mathrm{spec}(\phi)$ is not only well-ordered but in fact discrete. However, Theorem 7.2 will be enough for our purposes.

Theorem 7.3 (Existence of minpoints)**.** *Let ϕ any element in* $\mathrm{Aut}(\Gamma)$*. Then there exists $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ that has not jumped and such that*

$$\lambda_\phi(X) = \lambda(\phi).$$

Proof. Let $X_i \in \mathcal{O}(\Gamma)$ be a minimizing sequence for λ_ϕ. Without loss of generality we may assume that the sequence $\lambda_\phi(\Delta_{X_i})$ is monotone decreasing and Theorem 7.2 implies that it is eventually constant. Therefore X_i can be chosen in a fixed simplex Δ. Corollary 5.17 concludes. □

An interesting corollary of Theorem 7.3 is that we can characterize (global) jumps extending the first statement of Theorem 5.14 from a local to a global statement.

Theorem 7.4. *Fix $\phi \in \mathrm{Aut}(\Gamma)$. Let $X \in \mathcal{O}(\Gamma)$ and let $X_\infty \in \partial_\infty \Delta_X$ be obtained from X by collapsing a ϕ-invariant core graph A. Then X_∞ has not jumped if and only if*

$$\lambda(\phi|_A) \leq \lambda_\phi(X_\infty).$$

Proof. Suppose that X_∞ has not jumped. Then there is a simplex Δ where X_∞ has not jumped, and the claim follows from Theorem 5.14 because $\lambda_{\phi|_A} \leq \lambda_{\phi|_A}(\Delta_A)$.

On the other hand, suppose $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. By Theorem 7.3 there is a simplex in $\mathcal{O}(A)$ containing a minimizing sequence for $\phi|_A$. Let A_ε be an element in that simplex so that $\lambda_{\phi|_A}(A_\varepsilon) < \lambda(\phi|_A) + \varepsilon$, and let $f_A : A_\varepsilon \to A_\varepsilon$ be an optimal map representing $\phi|_A$. Note that A_ε and A may be not homeomorphic. Let $\widehat{X}$ be a Γ-graph obtained by inserting a copy of A_ε in X_∞. (We notice that since A_ε may be not homeomorphic to A, we can have $\Delta_{\widehat{X}} \neq \Delta_X$. We also notice that such $\Delta_{\widehat{X}}$ is not unique as we have plenty of freedom of attaching the edges of X_∞ to A_ε.) By Lemma 5.12, for any $\varepsilon > 0$ there is an element $X_\varepsilon \in \Delta_{\widehat{X}}$ and an optimal map $f_\varepsilon : X_\varepsilon \to X_\varepsilon$ representing ϕ so that $X_\varepsilon \to X_\infty$ and $\mathrm{Lip}(f_\varepsilon) \leq \lambda_\phi(X_\infty) + \varepsilon$, hence $\lambda_\phi(X_\varepsilon) \leq (X_\infty) + \varepsilon$. Thus X_∞ has no jump in $\Delta_{\widehat{X}}$, and therefore has not jumped. □

By Theorem 4.17 we know that minimal displaced points and train tracks coincide. But some care is needed here, as that theorem is stated for point of $\mathcal{O}(\Gamma)$, and not for points at infinity. In fact, given $\phi \in \mathrm{Aut}(\Gamma)$, $X \in \mathcal{O}(\Gamma)$ and $A \subset X$ a ϕ-invariant sub-graph, a priori it may happen that $\lambda(\phi)$ is different if we consider ϕ as an element of $\mathrm{Aut}(X)$ or of $\mathrm{Aut}(X/A)$. That is to say, we may have $X_\infty = X/A$ such that $\lambda_\phi(X_\infty) = \lambda(\phi)$ but X_∞ is not a train track point in $\mathcal{O}(X/A)$.

For instance, consider the case where $X = A \cup B$, with both A and B invariant. Suppose that $\lambda(\phi) = \lambda(\phi|_A) > \lambda(\phi|_B)$. Now suppose that $\lambda(\phi) = \lambda_\phi(X) = \lambda_{\phi|_A}(A) = \lambda_{\phi|_B}(B)$. Collapse A. Then the resulting point X_∞ is a min point for ϕ in $\overline{\mathcal{O}(\Gamma)}^\infty$ which has not jumped, but since $\lambda(\phi|_B) < \lambda(\phi)$, it is not a min point for ϕ on $\mathcal{O}(X/A)$.

We want to avoid such a pathology. Here we need to make a difference between $\lambda(\phi)$ computed in different spaces, so we will specify the space over which we take the infimum.

Lemma 7.5 (Existence of train tracks)**.** *Let $\phi \in \mathrm{Out}(\Gamma)$. Let $X_\infty \in \overline{\mathcal{O}(\Gamma)}^\infty$ be such that:*

- *There is $X \in \mathcal{O}(\Gamma)$ such that X_∞ is obtained from X by collapsing a (possibly empty) core sub-graph A in X;*
- *$\lambda_\phi(X_\infty) = \inf_{Y \in \mathcal{O}(\Gamma)} \lambda_\phi(Y)$;*
- *it has no jump in Δ_X.*

Suppose moreover that X_∞ maximizes the dimension of Δ_{X_∞} among the set of elements in $\overline{\mathcal{O}(\Gamma)}^\infty$ satisfying such conditions (such a set is not empty by Theorem 7.3). Then $\lambda_\phi(X_\infty) = \inf_{Y \in \mathcal{O}(X/A)} \lambda_\phi(Y)$. (Hence it is in $\mathrm{TT}(\phi) \subset \mathcal{O}(X/A)$.)

Proof. If A is empty this is an instance of Theorem 4.17. Otherwise, suppose X_∞ is not a train track point of $\mathcal{O}(X_\infty)$. We claim that near X_∞ there is a point $X'_\infty \in \mathcal{O}(X_\infty)$ such that $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty)$. Indeed, if X_∞ is not a local min point in $\Delta_{X_\infty} \subset \mathcal{O}(X_\infty)$, then we can find X'_∞ just near X_∞ in Δ_{X_∞}. Otherwise, by Lemma 4.23 there is a point X'_∞ obtained form X_∞ by folds directed by optimal maps (and such that $\dim(\Delta_{X'_\infty}) > \dim(\Delta_{X_\infty})$) such that $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty)$.

Let $\varepsilon = (\lambda_\phi(X_\infty) - \lambda_\phi(X'_\infty))/2$.

Since X_∞ has not jumped, by Theorem 7.4 we have $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. If $\lambda(\phi|_A) < \lambda_\phi(X_\infty)$, let $A' \in \mathcal{O}(A)$ be a point such that $\lambda_{\phi|_A}(A') < \lambda_\phi(X_\infty)$. Now Lemma 5.12 provides an element of $\mathcal{O}(\Gamma)$ which is displaced less or equal than $\max\{\lambda_{\phi|_A}(A'), \lambda_\phi(X'_\infty) + \varepsilon\}$, contradicting the fact that X_∞ is a minpoint for λ. Therefore $\lambda(\phi|_A) = \lambda(X_\infty)$.

By Theorem 7.3 there is $A_\infty \in \overline{\mathcal{O}(A)}^\infty$ such that $\lambda_{\phi|_A}(A_\infty) = \lambda(\phi|_A)$ and which has not jumped in $\mathcal{O}(A)$. Thus A_∞ is obtained, without jumps, from a point $A'_\infty \in \mathcal{O}(A)$ by collapsing a (possibly empty) invariant core sub-graph B. So $A_\infty \in \mathcal{O}(A'_\infty/B)$.

Let Y be a Γ-graph obtained by inserting a copy of A'_∞ in X'_∞. Let Y' be the graph obtained collapsing B. Y' belongs to the simplicial boundary of Δ_Y and, since A_∞ has no jump, then so does Y'. Now, observe that $Y' \in \mathcal{O}(Y/B)$ and A_∞ is a ϕ-invariant subgraph of Y' so that $Y'/A_\infty = X'_\infty$. Lemma 5.12 provides an element in $Y'_\infty \in \mathcal{O}(Y/B)$, in the same simplex of Y' which is displaced no more than $\lambda_{\phi|_A}(A_\infty)$ (because $\lambda_\phi(X'_\infty) < \lambda_\phi(X_\infty) = \lambda_{\phi|_A}(A_\infty)$). Now, Y'_∞ is a new minpoint for λ with $\dim(\Delta_{Y'_\infty}) > \dim(\Delta_{X_\infty})$ contradicting the maximality hypothesis on X_∞. It follows that X_∞ is a train track point in $\mathcal{O}(X_\infty)$ as desired. □

So we have seen that, even if non-jumping min-points are not necessarily train tracks, some of them are. Conversely, we see now non-jumping train tracks at the bordification are always min-point for λ_ϕ.

Lemma 7.6. *Let $\phi \in \mathrm{Aut}(\Gamma)$ and let $X \in \mathcal{O}(\Gamma)$. If there is k so that there is a constant $A > 0$ such that for any $n >> 1$*

$$Ak^n \leq \Lambda(X, \phi^n X)$$

then $k \leq \lambda(\phi)$.

Proof. This follows from the multiplicative triangular inequality. For any $Y \in \mathcal{O}(\Gamma)$ we have $\Lambda(Y, \phi^n Y) \leq \Lambda(Y, \phi Y)^n$. Define a constant $C = \Lambda(X, Y)\Lambda(Y, X)$ and notice that we also have $C = \Lambda(X, Y)\Lambda(\phi Y, \phi X)$. Then,

$$Ak^n \leq \Lambda(X, \phi^n X) \leq \Lambda(X, Y)\Lambda(Y, \phi^n Y)\Lambda(\phi^n Y, \phi^n X) \leq C\Lambda(Y, \phi Y)^n$$

whence, for any n

$$\left(\frac{k}{\Lambda(Y, \phi Y)}\right)^n \leq \frac{C}{A}.$$

This implies $k \leq \Lambda(Y, \phi(Y))$. By choosing a minimizing sequence of points Y_i we get $k \leq \lambda(\phi)$. □

Lemma 7.7. *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let $X_\infty \in \overline{\mathcal{O}(\Gamma)}$ which has not jumped. Suppose that there is a loop $\gamma \in X_\infty$ and $k > 0$ such that $L_{X_\infty}(\phi^n)(\gamma) \geq k^n L_{X_\infty}(\gamma)$. Then*

$$k \leq \lambda(\phi).$$

In particular, if X_∞ is a train track for ϕ as an element of $\mathrm{Aut}(X_\infty)$, then it is a minpoint for ϕ as an element of $\mathrm{Aut}(\Gamma)$.

Proof. Let $X \in \mathcal{O}(\Gamma)$ so that X_∞ is obtained from X by collapsing a core sub-graph $A \subset X$. Let X_ε be a point of X where $\mathrm{vol}(A) < \varepsilon$. Let γ be as in the hypothesis. For ε small enough we have $L_{X_\varepsilon}(\gamma) \leq 10 L_{X_\infty}(\gamma)$, and therefore

$$\Lambda(X_\varepsilon, \phi^n X_\varepsilon) \geq \frac{L_{X_\varepsilon}(\phi^n \gamma)}{L_{X_\varepsilon}(\gamma)} \geq \frac{L_{X_\infty}(\phi^n \gamma)}{10 L_{X_\infty}(\gamma)} \geq \frac{k^n L_{X_\infty}(\gamma)}{10 L_{X_\infty}(\gamma)} = \frac{k^n}{10}.$$

By Lemma 7.6 we have $\lambda(\phi) \geq k$.

For the second claim it suffice to choose let γ a legal candidate that realizes $\Lambda(X_\infty, \phi X_\infty)$. So $L_{X_\infty}(\phi^n(\gamma)) = \lambda_\phi(X_\infty)^n L_{X_\infty}(\gamma)$.

Hence $\lambda(\phi) \geq \lambda_\phi(X_\infty)$ and since X_∞ has not jumped $\lambda(\phi) \leq \lambda_\phi(X_\infty)$. □

We are now in position of extending the second claim of Theorem 5.14.

Corollary 7.8. *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let $X \in \mathcal{O}(\Gamma)$ and X_∞ be obtained from X by collapsing a ϕ-invariant core sub-graph A. Then*

$$\lambda(\phi|_A) \leq \lambda(\phi).$$

Moreover, if $\lambda(\phi|_A) = \lambda_\phi(X_\infty)$, then

$$\lambda(\phi) = \lambda(\phi|_A).$$

In particular X_∞ has not jumped if and only if

$$\lambda(\phi) \leq \lambda(X_\infty).$$

Proof. Let $\lambda = \lambda(\phi|_A)$. By Lemma 7.5 and Theorem 4.17, there is $\bar{A} \in \overline{\mathcal{O}(A)}^\infty$ which is a min-point for $\phi|_A$, which has not jumped in $\mathcal{O}(A)$, and which is a train track for $\phi|_A$ as an element of $\mathrm{Aut}(\bar{A})$. Let f_A be a train track map $f_A : \bar{A} \to \bar{A}$ representing $\phi|_A$. Therefore, there is a legal loop γ in $\bar{A}_{\max}$ whit legal images in $\bar{A}_{\max}$ and stretched exactly by λ. Let now $\widehat{X}$ be a metric Γ-graph obtained by inserting a copy of $\bar{A}$ in X_∞. Since $\overline{A}$ has not jumped in $\mathcal{O}(A)$, then $\widehat{X}$ has not jumped in $\mathcal{O}(\Gamma)$.

Let $f : \widehat{X} \to \widehat{X}$ be any PL-map representing ϕ so that $f|_A = f_A$. Therefore $f_A^n(\gamma)$ is immersed for any n and the length of $f_A^n(\gamma)$ is λ^n times the length of γ. It follows that $L_{\widehat{X}}((\phi^n)\gamma) = \lambda^n(L_{\widehat{X}}(\gamma))$.

By Lemma 7.7 $\lambda(\phi|_A) = \lambda \leq \lambda(\phi)$, and the first claim is proved. Moreover, if $\lambda(\phi|_A) = \lambda_\phi(X_\infty)$, then

$$\lambda(\phi) \leq \lambda(X_\infty) = \lambda(\phi|_A) = \lambda \leq \lambda(\phi)$$

and therefore all inequalities are equalities. Finally, if X has no jumped then $\lambda(X) \geq \lambda(\phi)$ just because this inequality is true by definition for points in $\mathcal{O}(\Gamma)$ and clearly passes to limits of non-jumping sequences, and the converse inequality follows from the second claim and Theorem 7.4. □

Note that Corollary 7.8 implies that *a posteriori* we can remove the non-jumping requirement from Theorem 7.3 and Lemma 7.5.

Corollary 7.9 (Min-points don't jump)**.** *Let ϕ any element in* $\mathrm{Aut}(\Gamma)$*. If $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ is such that $\lambda_\phi(X) = \lambda(\phi)$, then it has not jumped.*

Proof. This is a direct consequence of Corollary 7.8. □

We introduce the notion of train track at infinity.

Definition 7.10 (Train track at infinity)**.** Let $\phi \in \mathrm{Aut}(\Gamma)$. The set $\mathrm{TT}^\infty(\phi)$ is defined as the set of points $X \in \overline{\mathcal{O}(\Gamma)}^\infty$ such that X has not jumped, and X is a train track point for ϕ in $\mathcal{O}(X)$. (Hence $\lambda_\phi(X) = \lambda(\phi)$ by Lemma 7.7.)

Note that $\mathrm{TT}(\phi) \subset \mathrm{TT}^\infty(\phi)$. The main differences are that $\mathrm{TT}(\phi)$ may be empty (if ϕ is thin) while any ϕ has a train track in $\mathrm{TT}^\infty(\phi)$. On the other side, $\mathrm{TT}(\phi)$ coincides with the set of minimally displaced points, while $\mathrm{TT}^\infty(\phi)$ may be strictly contained in the set of minimally displaced points.

With this definition we can collect some of the above results in the following simple statement, which is a straightforward consequence of Theorems 4.17, 7.3 and Lemmas 7.5, 7.7.

Theorem 7.11. *For any $\phi \in \mathrm{Out}(\Gamma)$, $TT^\infty(\phi) \neq \emptyset$. For any $X \in \mathrm{TT}^\infty(\phi)$, $\lambda_\phi(X) = \lambda(\phi)$.*

The following corollary shows that if ϕ is reducible then there is a train track showing reducibility.

Corollary 7.12 (Detecting reducibility)**.** *Let $\phi \in \mathrm{Aut}(\Gamma)$ be reducible. Then there is $T \in \mathrm{TT}^\infty(\phi)$ such that either $T \in \partial_\infty \mathcal{O}(\Gamma)$ or there is an optimal map $f_T : T \to T$ representing ϕ such that there is a proper sub-graph of T which is f_T-invariant.*

Proof. Since ϕ is reducible there is $X \in \mathcal{O}(\Gamma)$, a PL-map $f : X \to X$ representing ϕ and a proper sub-graph $A \subset X$ such that $f(A) = A$. We can therefore collapse A and λ won't explode. By Theorem 7.11 there is a train track Z for ϕ in $\overline{\mathcal{O}(X/A)}^\infty$ and a train track Y for $\phi|_A$ in $\overline{\mathcal{O}(A)}^\infty$. If $\lambda_{\phi|_A}(Y) \leq \lambda_\phi(Z)$, then $Z \in \mathrm{TT}^\infty(\phi) \cap \partial_\infty \mathcal{O}(\Gamma)$ and we are done. Otherwise, since Z has not jumped (as a point of $\partial_\infty \mathcal{O}(X/A)$), we can regenerate it to a point $Z' \in \mathcal{O}(X/A)$ with $\lambda_\phi(Z') < \lambda_{\phi|_A}(Y)$. We now apply regeneration Lemma 5.12 to Y and Z'. If $Y \in \partial_\infty \mathcal{O}(A)$, then we get a train track for ϕ in $\partial_\infty \mathcal{O}(\Gamma)$. If $Y \in \mathcal{O}(A)$ we get a train track for ϕ in $\mathcal{O}(\Gamma)$ admitting Y as an invariant sub-graph. □

In fact, the proof of Corollary 7.12 proves more: that train tracks detect any invariant free factor.

Corollary 7.13 (Strong reformulation of Corollary 7.12)**.** *Let $\phi \in \mathrm{Aut}(\Gamma)$. Let X be a Γ-graph having a ϕ-invariant core sub-graph A. Then there is $Z \in \mathcal{O}(X/A)$ and $W \in \mathrm{Hor}(Z)$ such that the simplex Δ_W contains a minimizing sequence for λ. Moreover if $Y \in \mathcal{O}(A)$ is the graph used to regenerate W from Z, then the minimizing sequence can be chosen with* PL*-maps f_i such that $f_i(Y) = Y$ and $\mathrm{Lip}(f_i) \to \lambda(\phi)$.*

Proof. Follows from the proof of Corollary 7.12 (and Lemma 5.12). □

Finally, as in the case of irreducible automorphisms, the existence of train tracks gives the following fact.

Corollary 7.14. *For any $\phi \in \mathrm{Aut}(\Gamma)$ we have $\lambda(\phi^n) = \lambda(\phi)^n$.*

Proof. It follows from Theorem 7.11 and Lemma 4.14. □

8. Statement of main theorem and regeneration of paths in the bordification

We use Notation 2.10, that we recall here for the benefit of the reader.

- G will always mean a group with a splitting $\mathcal{G} : G = G_1 * \cdots * G_p * F_n$;
- $\Gamma = \sqcup \Gamma_i$ will always mean that Γ is a finite disjoint union of finite graphs of groups Γ_i, each with trivial edge-groups and non-trivial fundamental group $H_i = \pi_i(\Gamma_i)$, each H_i being equipped with the splitting given by the vertex-groups.

Also, we recall that we defined the rank of Γ in Definition 2.11. Finally, we recall the notation for λ (Definition 4.2). For any automorphism $\phi \in \mathrm{Out}(\Gamma)$ we define the function

$$\lambda_\phi : \mathcal{O}(\Gamma) \to \mathbb{R} \qquad \lambda_\phi(X) = \Lambda(X, \phi X)$$

If Δ is a simplex of $\mathcal{O}(\Gamma)$ we define

$$\lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X)$$

If there is no ambiguity we write simply λ instead of λ_ϕ. Finally, we set

$$\lambda(\phi) = \inf_{X \in \mathcal{O}(\Gamma)} \lambda_\phi(X)$$

We agree that we extend the function λ to points in $X_\infty \in \partial_\infty(\mathcal{O}(\Gamma))$ for which there is a sequence of points $X_i \in \mathcal{O}(\Gamma)$ such that $X_i \to X_\infty$ with $\lambda(X_i)$ bounded above, and we set $\lambda = \infty$ on other points. (Definition 5.10.)

Definition 8.1. Let $X, Y \in \overline{\mathcal{O}(\Gamma)}^\infty$. A *simplicial path* between X, Y is given by:

(1) A finite sequence of points $X = X_0, X_1, \ldots, X_k = Y$, called vertices, such that $\forall i = 1, \ldots, k$, there is a minimal simplex Δ_i such that $\Delta_{X_{i-1}}$ and Δ_{X_i} are both simplicial faces of Δ_i (we allow one of them or even both to coincide with Δ_i).
(2) Euclidean segments $\overline{X_{i-1}X_i} \subset \Delta_i$, called edges.

Definition 8.2. We say that a set χ is *connected by simplicial paths* if for any $x, y \in \chi$ there is a simplicial path between x and y which is entirely contained in χ.

Theorem 8.3 (Level sets are connected)**.** *Let $\phi \in \mathrm{Out}(\Gamma)$. For any $\varepsilon > 0$ the set*

$$\{X \in \mathcal{O}(\Gamma) : \lambda_\phi(X) \leq \lambda(\phi) + \varepsilon\}$$

is connected in $\mathcal{O}(\Gamma)$ by simplicial paths. The set

$$\{X \in \overline{\mathcal{O}(\Gamma)}^{\infty} : \lambda_\phi(X) = \lambda(\phi)\}$$

is connected by simplicial paths in $\overline{\mathcal{O}(\Gamma)}^{\infty}$.

The remaining goal of the paper is devoted to the proof of Theorem 8.3. The rough strategy is to prove the second claim and then prove that paths in the bordification can regenerate to paths in $\mathcal{O}(\Gamma)$ without increasing λ too much. The proof goes by induction on the rank of Γ (see definition below).

Remark 8.4. Theorem 8.3 is trivially true if $\mathrm{rank}(\Gamma) = 1$, because in that case either $\mathcal{O}(\Gamma)$ or $\mathbb{P}\mathcal{O}(\Gamma)$ is a single point.

Lemma 8.5 (Regeneration of segments)**.** *Fix $[\phi] \in \mathrm{Out}(\Gamma)$. Let $X_\infty, Y_\infty \in \overline{\mathcal{O}(\Gamma)}^{\infty}$ such that Δ_{Y_∞} is a (not necessarily proper) simplicial face of Δ_{X_∞}. Suppose that $\lambda(X_\infty) \geq \lambda(\phi)$. Then there is an open simplex Δ of $\mathcal{O}(\Gamma)$ such that for any $\varepsilon > 0$ there is $Y \in \mathrm{Hor}(Y_\infty) \cap \overline{\Delta}$ and $X \in \mathrm{Hor}(X_\infty) \cap \Delta$ such that*

$$\lambda_\phi(Y), \lambda_\phi(X) < \max\{\lambda_\phi(Y_\infty), \lambda_\phi(X_\infty)\} + \varepsilon.$$

Moreover, such inequality holds on the whole segments $\overline{XX_\infty}$ and $\overline{YY_\infty}$.

Proof. Let X_∞ be obtained by collapsing a ϕ-invariant core-subgraph A from a Γ-graph $\widehat{X}$. Since $\lambda_\phi(X_\infty) \geq \lambda(\phi)$, by Corollary 7.8 $\lambda(\phi|_A) \leq \lambda_\phi(X_\infty)$. By Theorem 7.3 there is a simplex in $\mathcal{O}(A)$ that contains a minimizing sequence for $\lambda(\phi|_A)$. Let A_ε be a point in that simplex such that $\lambda(A_\varepsilon) < \lambda(\phi|_A) + \varepsilon$. The required simplex Δ is obtained by inserting a copy of A_ε at the place of A in X_∞. We notice that such a Δ is not unique. By Lemma 5.12 there is a point $X \in \Delta \cap \mathrm{Hor}(X_\infty)$ such that $\lambda_\phi(X) \leq \lambda_\phi(X_\infty) + \varepsilon$.

Let's now see what happens to the points in $\overline{\Delta} \cap \mathrm{Hor}(Y_\infty)$. By hypothesis there is a ϕ-invariant $B \subseteq X_\infty$ such that as a graph (i.e. forgetting the metrics), Y_∞ is obtained from X_∞ by collapsing B. B has a pre-image in X still denoted by B. Let T be the forest $(A \cup B) \setminus \mathrm{core}(A \cup B)$. If $Y' = X/T$, as a graph, $Y_\infty = X/(A \cup B) = Y'/\mathrm{core}(A \cup B)$.

Thus the finitary face $\Delta_{Y'}$ of Δ obtained by the collapse of T intersects $\mathrm{Hor}(Y_\infty)$.

Let $f : X \to X$ be an optimal map representing ϕ. Since $\mathrm{core}(A \cup B)$ is ϕ-invariant, $f(\mathrm{core}(A \cup B)) \subset \mathrm{core}(A \cup B)$ up to homotopy. It follows that there is a PL-map $g : \mathrm{core}(A \cup B) \to \mathrm{core}(A \cup B)$ representing $\phi|_{A \cup B}$ such that $\mathrm{Lip}(g) \leq \lambda_\phi(X) \leq \lambda_\phi(X_\infty) + \varepsilon$. By Lemma 5.12 there is a point $Y \in \mathrm{Hor}(Y_\infty) \cap \Delta_{Y'}$ such that $\lambda_\phi(Y) \leq \max\{\lambda_\phi(Y_\infty) + \varepsilon, \mathrm{Lip}(g)\} \leq \max\{\lambda_\phi(Y_\infty) + \varepsilon, \lambda_\phi(X_\infty) + \varepsilon\}$. The last claim also follows by Lemma 5.12. □

Now we can plug in the inductive hypothesis in the proof of Theorem 8.3. Recall that if $X = T/S$ as graphs of groups, then we denote by $\pi : \mathrm{Hor}(X) \to \mathbb{P}\mathcal{O}(S)$ the projection that associates to a point in $\mathrm{Hor}(X)$ its collapsed part (see section 2.5).

Lemma 8.6 (Regeneration of horospheres)**.** *Suppose that Theorem 8.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$*. Let* $\phi \in \mathrm{Out}(\Gamma)$*. Let* $T \in \mathcal{O}(\Gamma)$ *be a* Γ*-graph having a proper* ϕ*-invariant core sub-graph* S*. Let* $X \in \partial_\infty \mathcal{O}(\Gamma)$ *be the graph obtained from* T *by collapsing* S*, and let* $A, B \in \mathrm{Hor}(X) \subset \mathcal{O}(\Gamma)$*. Let* m_A *and* m_B *be the supremum of* λ_ϕ *on the Euclidean segments* $\overline{AX}$ *and* $\overline{BX}$ *respectively. Then, for any* $\varepsilon > 0$ *there is a simplicial path* γ *between* A *and* B*, and in* $\mathrm{Hor}(X)$*, such that for any vertex* Z *of* γ *we have*

$$\lambda_\phi(Z) < \max\{m_A, m_B\} + \varepsilon.$$

Proof. Let $L = \max\{m_A, m_B\}$. Since S is ϕ-invariant, by Lemma 5.8 we have that $\lambda_\phi(X)$ is finite and by Lemma 5.12 both m_A and m_B are finite.

For any $Y \in \mathrm{Hor}(X)$, Theorem 3.7 implies $\lambda_\phi(\pi(Y)) \leq \lambda_\phi(Y)$ so

$$\lambda_\phi(\pi(A)) \leq \lambda_\phi(A) \quad \lambda_\phi(\pi(B)) \leq \lambda_\phi(B)$$

hence, $\lambda_\phi(\pi(A)), \lambda_\phi(\pi(B)) \leq L$. The rank of S is strictly smaller than $\mathrm{rank}(\Gamma)$ because it is a proper sub-graph of T. Hence Theorem 8.3 holds for $\mathcal{O}(S)$. So there is a finite simplicial path $(Y_i) \in \mathcal{O}(S)$ between $\pi(A)$ and $\pi(B)$ such that $\lambda_\phi(Y_i) < L + \varepsilon$. Then, there is a finite simplicial path in $\mathrm{Hor}(X)$ between A and B whose vertices are points $\widehat{T}_j$ such that for any j there is i such that $\pi(\widehat{T}_j) = Y_i$. By Lemma 5.12 there is a simplicial path in $\mathrm{Hor}(X)$ whose vertices are points $Z_j \in \Delta_{\widehat{T}_j}$ such that $\pi(Z_j) = \pi(\widehat{T}_j) = Y_i$ and $\lambda_\phi(Z_j) < L + \varepsilon$. □

We recall that we are using the notation of Definition 8.1.

Theorem 8.7 (Regeneration of paths)**.** *Suppose that Theorem 8.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$*. Let* $\phi \in \mathrm{Out}(\Gamma)$*. Let* $\gamma = (X_i)$ *be a simplicial path in* $\overline{\mathcal{O}(\Gamma)}^\infty$ *such that for every* i *either* $\Delta_{X_{i-1}}$ *is a simplicial face of* Δ_{X_i} *vice versa,*

Suppose that there is L *so that for any point* X_i *we have*

$$\lambda(\phi) \leq \lambda_\phi(X_i) \leq L.$$

Then, for any $\varepsilon > 0$ *there exists a simplicial path* η *in* $\mathcal{O}(\Gamma)$*, contained in the level set* $\lambda_\phi^{-1}(L+\varepsilon)$*, and such that each vertex of* η *belongs to the horosphere of some* X_j*.*

Proof. By Lemma 6.2 it suffices to define the vertices of the path η. By Lemma 8.5 For any i there are points $A_i, B_i \in \mathrm{Hor}(X_i)$ such that $\lambda_\phi(A_i), \lambda_\phi(B_i) \leq L + \varepsilon$ and such that B_i, A_{i+1} are in the same closed simplex of $\mathcal{O}(\Gamma)$. By Lemma 8.6 there is a simplicial path Y_{ij} between

A_i and B_i such that $Y_{ij} \in \mathrm{Hor}(X_i)$ and $\lambda_\phi(Y_{ij}) \leq L + \varepsilon$. The path η is now defined by the concatenation of such paths and the segments $\overline{B_i A_{i+1}}$. □

The proof of Theorem 8.3 now continues by an argument of peak reduction among simplicial paths connecting two points in the same level set. In next section we prove the results that will allow to reduce peaks.

9. Preparation to peak reduction

We keep Notation 2.10. We also recall that for $\phi \in \mathrm{Aut}(\Gamma)$ and a simplex $\Delta \in \overline{\mathcal{O}(\Gamma)}^\infty$ we are using the notation

$$\lambda(\Delta) = \lambda_\phi(\Delta) = \inf_{X \in \Delta} \lambda_\phi(X).$$

For the remaining of the section we fix $\phi \in \mathrm{Aut}(\Gamma)$. Recall that we are using the notation of Definition 8.1 for simplicial paths. In Theorem 8.7 we required that given two consecutive points X_{i-1}, X_i then one of $\Delta_{X_{i-1}}, \Delta_{X_i}$ is a face of the other. In general such a condition is easy to obtain by adding a middle point, but we need to do it in such a way to control the function λ, which is not in general continuous on $\overline{\mathcal{O}(\Gamma)}^\infty$.

We describe now a procedure for locally minimizing λ on simplicial path in $\mathcal{O}(\Gamma)$.

Let $(X_i)_{i=0}^k$ be a simplicial path such that:

- $X_2, \dots, X_{k-1} \in \mathcal{O}(\Gamma)$;
- If $X_0 \notin \mathcal{O}(\Gamma)$ then $X_0 \in \partial_\infty \Delta_{X_1}$ and has no jump in Δ_{X_1}.
- If $X_k \notin \mathcal{O}(\Gamma)$ then $X_k \in \partial_\infty \Delta_{X_{k-1}}$ and has no jump in $\Delta_{X_{k-1}}$.

Then, we define a new simplicial path by doing the following steps:

(1) For any i, if X_{i-1} and X_i are both proper faces of Δ_i, then we add to the path a point $\widehat{X}_i \in \Delta_i$.
(2) We renumber the sequence of vertices, still denoted by X_i. So now the sequence is $(X_i)_{i=0}^m$ for some $m \geq k$.
(3) We set $Y_0 = X_0$ and $Y_m = X_m$.
(4) For any any i, we chose $Y_i \in \overline{\Delta_{X_i}}^\infty$ so that $\lambda(Y_i) = \lambda(\Delta_{X_i}) = \lambda(\Delta_{Y_i})$, moreover we require that Y_i maximizes the dimension of Δ_{Y_i} among such points. Such a Y_i exists and has not jumped in Δ_{X_i} by Corollary 5.17.
(5) If it happens that $Y_i = Y_{i+1}$ for some i, then we identity such points and we renumber the sequence accordingly.

Definition 9.1. We say that the path (Y_i) as above is obtained by *minimizing* the path (X_i). A path in $\overline{\mathcal{O}(\Gamma)}^\infty$ is said *minimized* if it is the optimization of a simplicial path $(X_i)_{i=0}^k$ contained in $\mathcal{O}(\Gamma)$ except at most at its endpoints X_0, X_k, which have no jump in Δ_{X_0} and $\Delta_{X_{k-1}}$ respectively.

Note that if (Y_i) has the same endpoints of (X_i), and any other vertex minimizes λ in its simplex. In particular

$$\sup_i \lambda(Y_i) \leq \sup_i \lambda(X_i).$$

Lemma 9.2. *Let A, B two consecutive vertices of a minimized simplicial path Σ.*

- *Por any point P of $\overline{AB}$ we have $\lambda(P) \geq \lambda(\phi)$;*
- *if $\lambda(A) = \lambda(B)$, then λ is constant on the segment $\overline{AB}$;*
- *if $\lambda(A) > \lambda(B)$ then λ is continuous and strictly monotone near A.*

Proof. By definition of minimized path, there is a sequence $A_i \to A$ and $B_i \to B$, both without jump and in the same closed simplex $\overline{\Delta}$, and one of them is in the open simplex Δ of $\mathcal{O}(\Gamma)$. Without loss of generality we may assume $\Delta = \Delta_{B_i}$. Clearly $\lambda(A), \lambda(B) \geq \lambda(\Delta) \geq \lambda(\phi)$ because they did not jumped in Δ. If $B \in \Delta$ then $P \in \Delta$, and

$$\lambda(P) \geq \lambda(\Delta) \geq \lambda(\phi).$$

If $B \notin \Delta$, by Corollary 5.17 no point of $\overline{\Delta}^\infty$ has jumped in Δ, so P has not jumped and again $\lambda(P) \geq \lambda(\Delta) \geq \lambda(\phi)$.

Suppose now that $\lambda(A) = \lambda(B) = L$. Since B minimizes λ on Δ we have $L = \lambda(\Delta)$. By quasi convexity

$$\lambda(\Delta) \leq \lambda(P) \leq \lambda(B) = \lambda(\Delta).$$

The last claim follows directly from Lemma 6.2 and Lemma 6.3. □

Note that in particular, Lemma 9.2 implies that λ is bounded from below by $\lambda(\phi)$ on any minimized simplicial path.

Definition 9.3. A simplicial path Σ is said *L-calibrated* if λ is continuous on Σ and for any point P of Σ we have

$$\lambda(\phi) \leq \lambda(P) \leq L,$$

and if max-points minimize λ on their simplices. (That is to say, if X is such that $\lambda(X) = \max_P \in \Sigma \lambda(P)$, then $\lambda(X) = \lambda_{\Delta_X}$.)

Note that by Lemma 6.2 if A, B are consecutive vertices of a calibrated path such that $\lambda(A) > \lambda(B)$ then λ is strictly monotone near A.

Lemma 9.4. *Let Σ be a minimized simplicial path and let*

$$L = \max_{P \in \Sigma} \lambda(P).$$

Then there is a L-calibrated simplicial path obtained by adding some extra vertices to Σ.

Proof. By Lemma 9.2 we have to care only about continuity. Let A, B be two consecutive points of Σ. By Lemma 9.2 λ can be not continuous only at the endpoint with lower λ. Suppose $\lambda(A) > \lambda(B)$. By definition of minimized path there is a sequence $A_i \to A$ and $B_i \to B$, both without jump and in the same closed simplex $\overline{\Delta}$, and one of them is in the open simplex Δ of $\mathcal{O}(\Gamma)$. Moreover, since $\lambda(B) = \lambda(\Delta_{B_i})$, $\lambda(A) = \lambda(\Delta_{A_i})$ and one of them equals $\lambda(\Delta)$, $\lambda(B) < \lambda(A)$ forces $\Delta = \Delta_{B_i}$. If $B \in \Delta$ then λ is continuous at B because λ is continuous on any open simplex Δ of $\mathcal{O}(\Gamma)$. If $B \notin \Delta$, then B has not jumped in Δ. Therefore (by Theorem 5.14 and Lemma 5.12), there is a point $\widehat{B} \in \Delta$ such that

- $\lambda(\widehat{B}) < \lambda(A)$;
- λ is continuous on the segment $\overline{\widehat{B}B}$.

by Lemma 6.3 λ is continuous on the segment $\overline{A\widehat{B}}$. Since we did not modified Σ at its max-points, also the second condition of calibration is assured because it is already satisfied for minimized paths. □

We prove now a (technical) fact that can be informally phrased as follows[8]:

> Given $X \in \overline{\mathcal{O}(\Gamma)}^{\infty}$ and $f : X \to X$ an optimal map representing ϕ, if Y is sufficiently close to X for the Euclidean metric, then any fold in X directed by f closely reads in Y.

Theorem 9.5. *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$. Suppose that Δ_X is a simplicial face of Δ_Y. Thus as graphs, Y is obtained by collapsing a sub-graph A. Suppose that* core(A) *is ϕ-invariant. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing A and by linearly rescaling the edges in $Y \setminus A$.*

Let $f : X \to X$ be an optimal map representing ϕ. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \leq t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing ϕ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. The proof of this theorem relies on accurate (but boring) estimates. For the happiness of the reader we postpone the proof to the appendix. □

Corollary 9.6. *In the hypotheses, and with notation of Theorem 9.5, let τ be an f-illegal turn of X and let Δ^τ be the simplex obtained from Δ_X by fold a little that turn. Then, for any $\varepsilon > 0$, there is t_ε such that $\forall t < t_\varepsilon$, there is a finite simplicial path Σ_t in $\mathcal{O}(Y)$ with vertices*

[8]We recall that by definition $\overline{\mathcal{O}(\Gamma)}^{\infty} = \overline{\mathcal{O}(\Gamma)}$ and that the symbol ∞ is just to put emphasis on the fact that we are considering the simplicial bordification of the outer space obtained by adding all simplices at infinity.

$Z_0^t = Y_t, Z_1^t, \dots, Z_m^t$ such that $\Delta_{Z_i^t}$ has Δ_X as a simplicial face for $i \neq m$, $\Delta_{Z_m^t}$ has Δ^τ as a simplicial face, and such that for any point Z of Σ_t we have

$$\lambda(X) - \varepsilon < \lambda(Z) \leq \lambda(Y_t).$$

Moreover, for $s \in [0,t]$, $s \mapsto Z_i^s$ parametrizes the segment from X to Z_i^t, and Z_m^s that from X^τ to Z_m^t.

Proof. For this proof we will work entirely with trees. So Y will denote a Γ-trees, A an equivariant family of sub-trees, and so on.

We denote by A_t the metric copy of A in Y_t. By hypothesis there are two different segments α_τ, β_τ incident at the same vertex v in X such that f overlaps α_τ and β_τ. If $v \notin \sigma_t(A_t)$ then, for any small enough ε and $t < t_\varepsilon$, also g_t must overlap $\alpha = \sigma_t^{-1}(\alpha_\tau)$ and $\beta = \sigma_t^{-1}(\beta_\tau)$, and the claim follows by (equivariantly) performing the corresponding simple fold directed by g_t. The inequality "$\leq \lambda(Y_t)$" follows because the fold is directed by an optimal map, the inequality "$> \lambda(X) - \varepsilon$" follows by lower semicontinuity of λ.

Otherwise, α and β are segments incident to the same component of A_t. If α and β are incident to the same point, then we proceed as above, so we can suppose that they are incident to different points of A.

For small enough ε and $t < t_\varepsilon$ we have that g_t overlaps some open sub-segments of α and β. Let $a \in \alpha$ and $b \in \beta$ such that $g_t(a) = g_t(b)$ and such that a is the closest possible to A.

Let γ be the shortest path from α and β in A_t. It turns out that γ is a simple simplicial path. On γ we put an extra simplicial structure given by the pull-back via g_t: we declare new vertices of γ the points whose g_t-image is a vertex of Y_t. $g_t(\gamma)$ is a tree because Y_t is. Moreover, since $g_t(a) = g_t(b)$, the restriction of g_t to γ cannot be injective. In particular, if $x \in \gamma$ is a point such that $d_{Y_t}(g_t(x), g_t(a))$ is maximal, then x is a vertex of γ, and the two sub-segments of γ incident to x are completely overlapped.

Let Z_1^t be the tree obtained by equivariantly identify such segments. Clearly, g_t induces a map $g_t^1 : Z_1^t \to Z_1^t$. Such map is continuous and not necessarily PL. However,

$$\mathrm{Lip}(g_t^1) \leq \mathrm{Lip}(g_t)$$

and $\mathrm{PL}(g_t^1)$ still represents ϕ. Since $\mathrm{Lip}(\mathrm{PL}(g_t^1)) \leq \mathrm{Lip}(g_t^1)$ he have

$$\lambda(Z_1^t) \leq \lambda(Y_t).$$

Note also that $\Delta_{Z_1^t}$ has Δ_X as a simplicial face because our identification occurred in A_t. Also, since Y_t parametrizes the segment from X to Y, as t varies Z_1^t parametrizes the segment from X to Z_1^t.

Note that a priori we may have $\Delta_{Z_1^t} = \Delta_Y$, but in any case $\Delta_{Z_1^t}$ is either a (non necessarily proper) simplicial face of Δ_Y or vice versa.

In Z_1^t we have a simple path γ_1 resulting from γ by the cancellation of the two identified segments at x. By construction g_t^1 is simplicial and not injective on γ_1. Therefore we can iterate the above procedure and define points Z_i^t with

$$\lambda(Z_i^t) \leq \mathrm{Lip}(g_t) = \lambda(Y_t)$$

and such that $\Delta_{Z_i^t}$ has Δ_X as a simplicial face. Moreover either $\Delta_{Z_i^t}$ has $\Delta_{Z_{i-1}t}$ as a simplicial face or vice versa. Since γ has a finite number of vertices, we must stop, and we do when γ_i is a single point. At this stage, α and β are incident to the same point and we are reduced to the initial case. Note that any $Z_i^t \to X$ as $t \to 0$, thus so does any point in segment from Z_i^t to Z_{i+1}^t. Therefore by lower semicontinuity of λ for any $\varepsilon > 0$, since we have finitely many points, for sufficiently small t we have that for any i

$$\lambda(X) - \varepsilon < \lambda(Z_i^t)$$

and the same inequality holds for points in the segments from Z_i^t to Z_{i+1}^t. □

Corollary 9.7. *In the hypothesis of Theorem 9.5, suppose that X is an exit point for Δ_X[9], and let X_E as Definition 4.22. Then, for any $\varepsilon > 0$, there is t_ε such that for any $t < t_\varepsilon$, there is a finite simplicial path Σ_t in $\mathcal{O}(Y)$ with vertices $Y_t = Z_0^t, Z_1^t, \ldots, Z_k^t$ such that $\Delta_{Z_i^t}$ has Δ_X as a simplicial face for $i \neq k$, $\Delta_{Z_k^t}$ has Δ_{X_E} as a simplicial face, and such that for any point Z of Σ_t we have*

$$\lambda(X) - \varepsilon < \lambda(Z) \leq \lambda(Y_t).$$

Moreover, for $s \in [0,t]$, $s \mapsto Z_k^s$ parametrizes the segment from from X_E to Z_k^t.

Proof. It follows by recursively apply Corollary 9.6. The uniform estimate on t follows because the path from X to X_E is finite. □

10. The end of the proof of Theorem 8.3: peak reduction on simplicial paths

We fix Γ as in Notation 2.10 and $\phi \in \mathrm{Aut}(\Gamma)$. Let $\lambda = \lambda_\phi$.

We will prove that for any $\varepsilon > 0$, the set

$$\{X \in \overline{\mathcal{O}(\Gamma)}^\infty : \lambda(\phi) \leq \lambda_\phi(X) \leq \lambda(\phi) + \varepsilon\}$$

is connected by $\lambda(\phi) + \varepsilon$-calibrated simplicial paths. This in particular gives the second claim of Theorem 8.3.

Moreover, if Σ is calibrated, then by possibly adding some extra vertices to Σ we obtain a path in the same level set that satisfies the hypotheses of Theorem 8.7 and therefore can be regenerated to $\mathcal{O}(\Gamma)$. Therefore, this proves also the first claim of Theorem 8.3.

[9]See Definition 4.22

From now on we fix $A, B \in \overline{\mathcal{O}(\Gamma)}$ such that $\lambda(A), \lambda(B) \geq \lambda(\phi)$. Let $L \geq \max\{\lambda(A), \lambda(B)\}$.

Let $\Sigma_L(A, B)$ be the set of L-calibrated simplicial paths from A to B.

Lemma 10.1. *For some L, $\Sigma_L(A, B) \neq \emptyset$.*

Proof. Since $\lambda(A), \lambda(B) \geq \lambda(\phi)$, they have not jumped. Let $A' \in \mathrm{Hor}(A)$ and $B' \in \mathrm{Hor}(B)$. Since $A', B' \in \mathcal{O}(\Gamma)$, which is connected, there is a simplicial path in $\mathcal{O}(\Gamma)$ between A', B'. After minimizing such path, by Lemma 9.4 we obtain an element of Σ_L (where the L is the maximum displacement along such a path). □

Definition 10.2. For any simplicial path $\Sigma = (X_i)$ we define $\max(\Sigma) = \max_{X_i} \lambda(X_i)$, and we say that X_i is a *peak* if $\lambda(X_i) = \max(\Sigma)$. A pair of two consecutive peaks X_{i-1}, X_i is called a *flat peak*. A peak is *strict* if it is not part of a flat peak.

Let $\Sigma_0 = (X_i) \in \Sigma(A, B)$ such that among all elements $\sigma \in \Sigma$ it minimizes, in order

(1) $\max(\Sigma)$
(2) the number peaks;
(3) the number of flat peaks.

Lemma 10.3. *Such a Σ_0 exists.*

Proof. By Theorem 7.2 we are minimizing over a well-ordered set. □

Note that if X is a strict peak then λ it is strictly monotone on both sides of X. (By Lemma 6.2.)

Once again, we need the inductive hypothesis.

Lemma 10.4. *Suppose that Theorem 8.3 is true in any rank less than* $\mathrm{rank}(\Gamma)$. *Then Σ_0 has no strict peaks in its interior.*

Proof. Suppose that $\lambda(X_{i-1}) < \lambda(X_i) > \lambda(X_{i+1})$. In particular we have $\lambda(\phi) < \lambda(X_i)$; a strict inequality. By calibration Δ_{X_i} minimizes λ in its simplex, hence Δ_{X_i} is a proper face of both Δ_i and Δ_{i+1}. Thus for any Y and Z, respectively in $\overline{X_{i-1}X_i}$ and $\overline{X_iX_{i-1}}$ we have

$$\lambda(\phi) < \lambda(Y), \lambda(Z) < \lambda(X_i).$$

We set $X = X_i$. If C is the collapsed part of X, then by Corollary 7.8

$$\lambda(\phi|_C) < \lambda(X).$$

As this is an open condition, it is preserved in an open neighbourhood U of X in $\mathcal{O}(X)$.[10]

Since X is not a ϕ-minimally displaced point, by Lemma 7.7 $X \notin \mathrm{TT}(\phi) \subset \mathcal{O}(X)$. By Lemma 4.23, X is an exit point. Let X_E as in Definition 4.22.

[10]Note that $\mathcal{O}(X)$ may be different from $\mathcal{O}(\Gamma)$.

Now we invoke Corollary 9.7. With the terminology of Corollary 9.7, let Y_t parametrize the segment from X to Y. By Corollary 9.7 there exists a simplicial path Σ_Y in $\mathcal{O}(Y)$ connecting Y_t to a point in $Y_E \in \mathrm{Hor}(X_E)$ (as a subset of $\mathcal{O}(Y)$).

Moreover, Corollary 9.6 gives the estimate

$$\lambda(X) - \varepsilon < \lambda(P) \leq \lambda(Y_t).$$

for any point of Σ_Y. In particular, if ε is mall enough we have

$$\lambda(\phi) < \lambda(P) < \lambda(X)$$

Strict inequalities. Similarly, there is a simplicial path Σ_Z connecting Z_t to a point $Z_E \in \mathrm{Hor}(X_E)$ with the same estimate above.

Since Σ_Y is in $O(Y)$ and λ is continuous on $\mathcal{O}(Y)$, then λ is continuous on Σ_Y. The same for Σ_Z.

Let $\widetilde{Y}_E, \widetilde{Z}_E \in \mathcal{O}(\Gamma)$ be points respectively in $\mathrm{Hor}(Y_E) \cap \mathrm{Hor}(X_E)$ and $\mathrm{Hor}(Z_E) \cap \mathrm{Hor}(X)$ such that $\lambda(\widetilde{Y}_E), \lambda(\widetilde{Z}_E) \leq \lambda(X)$ and such that λ is continuous on $\overline{Y_E\widetilde{Y}_E}$ and $\overline{Z_E\widetilde{Z}_E}$ (such points exist by Lemma 5.12).

By Lemma 8.6 there is a path Θ_E in $\mathcal{O}(\Gamma)$ and in the level set $\{\lambda(P) < \lambda(X)\}$ connecting $\widetilde{Y}_E$ to $\widetilde{Z}_E$. We add Y_E and Z_E to such points, we minimize and then apply Lemma 9.4. The result is a calibrated path θ_E connecting Y_t to Z_t in the level set $\{\lambda < X\}$.

The path obtained by following Σ till Y_t, then θ_E till Z_t and then Σ again, has either a lower maximum than Σ or one peak less than Σ. □

Lemma 10.5. *Σ_0 has no flat peaks unless λ is constant on Σ and $\lambda(\Sigma) = \lambda(\phi)$.*

Proof. If λ is not constantly $\lambda(\phi)$ on Σ, in particular λ is strictly bigger than $\lambda(\phi)$ on peaks.

Suppose that there is Y, X two consecutive vertices of Σ_0 with $\lambda(X) = \lambda(Y) = \max(\Sigma) > \lambda(\phi)$. The idea is to find a third point Z to add between Y and X in order to destroy the flat peak.

If there is a point Z in the interior of the segment YX, with $\lambda(\phi) \leq \lambda(Z) < \lambda(X) = \lambda(Y)$, then we add it.

Otherwise, λ is constant on $\overline{XY}$. Let W be a point in the interior of the segment $\overline{XY}$. If W is not a local minimum for λ in Δ_W, then near W we find Z with the above properties. We add it.

If W is a local minimum for λ in Δ_W then, by Lemma 4.23, near W in $\mathcal{O}(W)$ there is a point Z with the above properties and such that Δ_W is a finitary face of Δ_Z in $\mathcal{O}(W)$. We add Z.

Since λ is continuous on closed simplices of $\mathcal{O}(W)$, then by adding Z in each of the above case we obtain a new L-calibrated path Σ_1 which has the same number of peaks and exactly one flat peak less then Σ_0. This is impossible by the minimality of Σ_0. □

To finish the proof of Theorem 8.3, simply observe that we have shown that we can connect any two points in $\{X \in \overline{\mathcal{O}(\Gamma)}^{\infty} : \lambda_\phi(X) =$

$\lambda(\phi)\}$ by a calibrated simplicial path with no peaks, either strict or flat. This immediately implies that the displacement is constant along the path. □

11. Applications

In this section we show how the connectedness of the level sets gives a solution to the conjugacy problem.

We start with some technical results. Recall that a point, X, of CV_n is called ϵ-thin if there is a homotopically non-trivial loop in X of length at most ϵ. Conversely, X is called ϵ-thick if it is not ϵ-thin.

Proposition 11.1 ([3], Proposition 10)**.** *Let $X \in CV_n$ (that is, X is a marked metric graph) and $f : X \to X$ a PL-map representing some automorphism of F_n. Let $\lambda = Lip(f)$, let N equal maximal chain of topological subgraphs of any graph in CV_n (this is clearly a finite number) and let μ be any real number greater than λ. Then if X is $1/((3n-3)\mu^{(N+1)})$-thin, the automorphism represented by f is reducible. For instance, one can take $N = 3n - 3$.*

Definition 11.2. Let $X \in CV_n$. Then we call R an adjacent uniform rose if it obtained by collapsing a maximal tree in X and then rescaling so that all edges in R have the same length (that is, $1/n$, as we will work with volume 1).

Proposition 11.3. *Let $X \in CV_n$ be a point which is ϵ-thick and let R be any adjacent uniform rose (both of volume 1). Then, $\Lambda(X, R) \leq 1/\epsilon$ and $\Lambda(R, X) \leq n$.*

Proof. By Theorem 3.7, we can look at candidates that realise the stretching factor. Since, topologically, one passes from X to R by collapsing a maximal tree, we get that a candidate in X, when mapped to R, crosses every edge at most twice. In fact the candidate crosses every edge of R at most once in the case of an embedded simple loop or an infinity loop. This gives the first inequality, on taking into account that X is ϵ-thick and that barbells have length at least 2ϵ.

For the second inequality note that a embedded loop in R_0 is a edge and has length $1/n$ and lifts to an embedded loop in X, of length at most 1. An infinity loop in R_0 consists of two distinct edges, has length $2/n$ and lifts to a loop in X which goes through every edge at most twice. (Barbells are not present in R_0). □

Corollary 11.4. *Let $X \in CV_n$ be ϵ-thick and let R be an adjacent uniform rose. Consider $\Phi \in Out(F_n)$. Then $\Lambda(R, \phi R) \leq \frac{n}{\epsilon}\Lambda(X, \phi X)$.*

Proposition 11.5. *Let R, R_∞ be two points in CV_n which are both uniform roses (graphs with exactly one vertex and so that every edge has the same length). Let $\phi \in Out(F_n)$ be irreducible and suppose that μ is any real number greater than* $\max(\Lambda(R, \phi R), \Lambda(R_\infty, \phi R_\infty))$.

Then there exist $R_0 = R, R_1, R_2, \ldots, R_k = R_\infty$, which are all uniform roses in CV_n such that:

- *For each i, there exists a simplex Δ_i such that Δ_{R_i} is a rose face of both Δ_i and Δ_{i+1}.*
- *$\Lambda(R_i, \phi R_i) \leq \frac{n}{\epsilon}\mu$, where $\epsilon = 1/((3n-3)\mu^{(N+1)})$.*

Proof. This follows from Theorem 8.3, using Definition 8.1, since each pair Δ_i and Δ_{i+1} have a (at least one) common rose face; just take any uniform adjacent rose in any common rose face. The remaining point follows from Corollary 11.4 and Proposition 11.1. □

Proof of Theorem 1.4: We clearly have an algorithm which terminates, and it is apparent that if $\psi \in S_\phi$ then these automorphisms are conjugate. It remains to show the converse; that if they are conjugate, then $\psi \in S_\phi$.

Let R be the uniform rose corresponding to the basis B. If ψ were conjugate to ϕ, then there would be a conjugator, some $\tau \in Out(F_n)$ such that $\psi = \tau^{-1}\phi\tau$. Let $R_\infty = \tau R$. Now use Proposition 11.5 to find a sequence $R = R_0, R_1, \ldots, R_k = R_\infty$, such that each consecutive pair are incident to a common simplex and $\Lambda(R_i, \phi R_i) \leq n(3n-3)\mu^{3n-1} = K$.

We let ζ_i be an automorphism which sends R_i to R_{i+1}; the fact that these roses are both incident to a common simplex implies that each ζ_i is a CMT automorphism. Inductively, we may define, $\tau_i = \zeta_0 \ldots \zeta_{i-1}$, and note that $\tau_i R = R_i$. We make these choices so that $\tau = \tau_k$. (This possible since regardless of the choices made, we always have that $\tau_k^{-1}\tau$ fixes R and is therefore a CMT automorphism, therefore by possibly adding a single repetition of roses at the start we may assume that $\tau = \tau_k$.)

Now let $\phi_i = \tau_i^{-1}\phi\tau_i$.

Since $\phi_{i+1} = \zeta_i^{-1}\phi_i\zeta$, to finish the proof we just need that $||\phi_i|| \leq K$. This follows since,

$$\Lambda(R_i, \phi R_i) = \Lambda(\tau_i R, \phi\tau_i R) = \Lambda(R, \phi_i R) = ||\phi_i||_B.$$

□

We conclude the paper by proving Theorem 1.8. First a lemma,

Lemma 11.6. *Let X be a core graph and f a homotopy equivalence on X, having a proper, homotopically non-trivial subgraph X_0 such that $f(X_0) = X_0$. Then there is a maximal tree, T, such that the automorphism induced by f on the rose X/T is visibly reducible.*

Proof. Choose X_0 to be minimal. Therefore it will have components, $X_1, \ldots, X_k$ such that $f(X_i) = X_{i+1}$ with subscripts taken modulo k. Take a maximal tree for each X_i and extend this to a maximal tree, T, for X. It is then clear that if we take B_i to be the set of edges in X/T coming from X_i, that ψ will be visibly reducible as witnessed

by $B_1, \ldots, B_k$. (Note each subgroups generated by each B_i are only permuted/preserved up to conjugacy, since the X_i are disjoint and so one cannot choose a common basepoint). □

Proof of Theorem 1.8: We proceed much as in the proof of Theorem 1.4, but here we do not know that the points in CV_n we encounter will remain uniformly thick.

The algorithm clearly terminates, and if there is a ψ in S^+ which is visibly reducible, then ϕ is reducible. It remains, therefore, to show that if ϕ is reducible, then there is some $\psi \in S^+$ which is visibly reducible.

Let R be the uniform rose corresponding to the basis B. By Corollary 7.13, there exists an $X \in CV_n$ with a core invariant subgraph and such that $\Lambda(X, \phi(X)) \leq \mu$.

By Theorem 8.3, there exist points, $X_0 = R, X_1, \ldots, X_k = X$, such that $\Lambda(X, \phi(X_i)) \leq \mu$. Choose the maximal index, M, such that $X_0, X_1, \ldots, X_M$ are all ϵ-thick, where $\epsilon = 1/((3n-3)\mu^{(N+1)})$ as in Lemma 11.1. Now for each $i \leq M$, choose an R_i which is a adjacent uniform rose to both X_i and X_{i+1} (choose $R_0 = R$ and if $M = k$, let R_k be any uniform rose adjacent to X_k).

If $M = k$, we set $R_{M+1} = R_M$. Otherwise, by Lemma 11.1, we have that X_{M+1} has an optimal PL-representative for ϕ which admits an invariant subgraph. So by Lemma 11.6, we may find an adjacent uniform rose face, R_{M+1} so that the representative, ψ, of ϕ at R_{M+1} is visibly reducible.

As above, we $\tau \in Out(F_n)$ such that $\psi = \tau^{-1}\phi\tau$. Then let ζ_i be an automorphism which sends R_i to R_{i+1}; each ζ_i is a CMT automorphism. Inductively, we may define, $\tau_i = \zeta_0 \ldots \zeta_{i-1}$, and note that $\tau_i R = R_i$. We make these choices so that $\tau = \tau_{M+1}$.

Now let $\phi_i = \tau_i^{-1}\phi\tau_i$, so that $\phi_0 = \phi$ and $\phi_{M+1} = \psi$. Since each $X_0, \ldots, X_M$ is ϵ-thick we get, by Corollary 11.4 that each $\phi_i \in S_i$ for $i \leq M$. Hence $\psi \in S^+$ and is visibly reducible. □

12. Appendix: proof of Theorem 9.5

In this section we give the proof of Theorem 9.5, which we recall

Theorem (Theorem 9.5)**.** *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$. Suppose that Δ_X is a simplicial face of Δ_Y. Thus as graphs, Y is obtained by collapsing a sub-graph A. Suppose that* core(A) *is ϕ-invariant. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing A and by linearly rescaling the edges in $Y \setminus A$.*

Let $f : X \to X$ be an optimal map representing ϕ. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \leq t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing ϕ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. We split the proof in two sub-cases. First when A is itself a core-graph, and then the case when core(A) is empty. Clearly the disjoint union of the two cases implies the mixed case.

We will work at once with graphs and trees, by using the usual $\widetilde{\ }$ notation: if X is a graph, $\widetilde{X}$ is its universal covering an for any object o (a point, a sub-set, a map, ...) $\tilde{o}$ is one of its lift to the universal coverings. Vice versa, given an object $\tilde{o}$, we understand that $\tilde{o}$ descends (by equivariance if for example o is a map) to an object o at the level of graphs.

Lemma 12.1. *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$. Suppose that as graphs of groups, X is obtained from Y by collapsing a ϕ-invariant core sub-graph $A = \sqcup A_i$. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing A and by linearly rescaling the edges in $Y \setminus A$.*

Let $f : X \to X$ be an optimal map representing ϕ. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \leq t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing ϕ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. We begin by fixing some notation. Firs of all, we will use the symbol λ to denote any displacement functions of ϕ (i.e. $\lambda_\phi, \lambda_{\phi|_A}, \dots$) If x is a point in a metric space, we denote by $B_r(x)$ the open metric ball centered at x and radius r. For any i, we denote by v_i the non-free vertex of X obtained by collapsing A_i. For any t we denote by A^t the metric copy of A in Y_t. Note that A is uniformly collapsed in Y_t, that is to say, $[A^t] \in \mathbb{P}\mathcal{O}(A)$ is the same element for any $0 < t \leq 1$, and we have $\mathrm{vol}(A^t) = t\,\mathrm{vol}(A^1)$.

By lower semicontinuity of λ (Theorem 5.8) we have that

$$\forall \varepsilon_0 > 0 \exists t_{\varepsilon_0} > 0 \text{ such that } \forall t < t_{\varepsilon_0} \text{ we have } \lambda(Y_t) > \frac{\lambda(X)}{1+\varepsilon_0}. \tag{2}$$

A priori f may collapse some edge, in any case $\forall \varepsilon_1 > 0 \exists f_1 : X \to X$ a PL- map representing ϕ such that f_1 does not collapse any edge, and

$$(3)\quad d_\infty(f, f_1) < \varepsilon_1 \quad \text{and} \quad \mathrm{Lip}(f_1) < \mathrm{Lip}(f)(1+\varepsilon_1) = \lambda(X)(1+\varepsilon_1).$$

Moreover $\exists 0 < \rho_0 = \rho_0(X, f_1)$ such that $\forall \rho < \rho_0$

- $B_\rho(x)$ is star-shaped for any $x \in X$ (i.e. it contains at most one vertex);
- for any i, each connected component of $f_1^{-1}(B_\rho(v_i))$ is star-shaped and contains exactly one pre-image of v_i;
- for any i, j the connected components of $f_1^{-1}(B_\rho(v_i))$ and those of $f_1^{-1}(B_\rho(v_j))$ are pairwise disjoint.

We fix an optimal map $\varphi : A^1 \to A^1$ representing $\phi|_A$. Since $[A^t] \in \mathbb{P}\mathcal{O}(A)$ does not depend on t, $\varphi : A^t \to A^t$ is an optimal map for any $t \in (0, 1]$ and the Lipschitz constant does not change. Clearly (by Sausage Lemma 3.7)

$$(4)\qquad \mathrm{Lip}(\varphi) \le \lambda(Y_t) \quad \text{for any } t.$$

The natural option is to define g_t by using $\sigma_t^{-1} \circ f_1 \circ \sigma_t$. Hence, we need to deal with places where σ_t^{-1} is not defined. First we fix a lift $\widetilde{\varphi}$ of ϕ.

Each germ of edge α at v_i in X corresponds to a germ $\alpha_Y (= \sigma_t^{-1}(\alpha))$ in Y incident to A_i at a point that we denote by p_α. For any such α we **choose** a lift $\widetilde{\alpha}$, that corresponds to a germ $\widetilde{\alpha}_Y$ incident to $\widetilde{p}_\alpha \in \widetilde{A}_i$. (See Figure 4.)

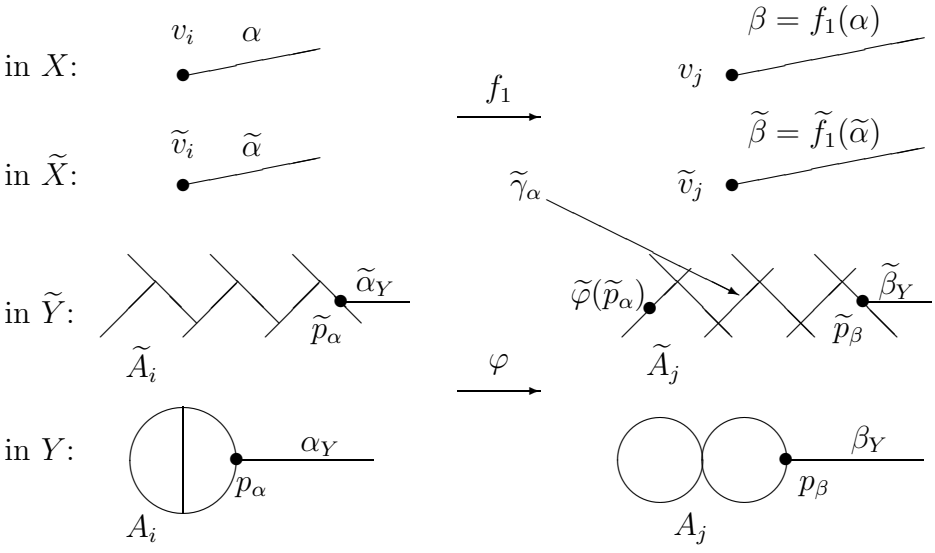

FIGURE 4. How to choose the paths $\widetilde{\gamma}_\alpha$

Suppose $f_1(v_i) = v_j$, and let $\beta = f_1(\alpha)$. Then $\widetilde{\beta}$ is a germ at $\widetilde{v}_j$ and corresponds to a germ $\widetilde{\beta}_Y$ incident to $\widetilde{A}_j$ at a point $\widetilde{p}_\beta$.

Let $\widetilde{\gamma}_\alpha$ be the unique path in $\widetilde{A}_j$ connecting $\widetilde{\varphi}(p_\alpha)$ to $\widetilde{p}_\beta$.

Remark 12.2. We choose a path $\widetilde{\gamma}_\alpha$ for any germ α in X, which is a finite graph. Therefore we have only finitely many such $\widetilde{\gamma}_\alpha$'s. We can then complete that family of paths by equivariance.

Now we do a similar construction for other pre-images of the v_i's. For any $x \in X$ such that $f_1(x) = v_i$ for some i, but $x \notin \{v_j\}$, we choose a base-point $\widetilde{x}_i \in \widetilde{A}_i$. Any germ of edge α at x correspond to an edge α_Y is Y (note that x is not necessarily a vertex of X). For any such α we **choose** a lift $\widetilde{\alpha}$. Since f_1 does not collapse edges, $\widetilde{f_1}(\widetilde{\alpha})$ is a germ of edge $\widetilde{\beta}$ at $\widetilde{v}_i$, and corresponds to a germ $\widetilde{\beta}_Y$ at $\widetilde{A}_i$ in $\widetilde{Y}$. Let $\widetilde{\gamma}_\alpha$ be the unique path in $\widetilde{A}_i$ connecting $\widetilde{x}_i$ and $\widetilde{\beta}_Y$.

Remark 12.3. As above we choose only finitely many such $\widetilde{\gamma}_\alpha$'s and we complete the choices equivariantly.

Note that, as germs, $\alpha_Y = \sigma_t^{-1}(\alpha)$ and $\beta_Y = \sigma_t^{-1}(\beta) = \sigma_t^{-1}(f_1(\alpha))$. Now we have a path $\gamma_\alpha \subset A$ for any pre-image of germs at the v_i's, chosen independently on t. Let $t \in (0, 1]$. We define a map

$$g : Y_t \to Y_t$$

representing ϕ as follows:

- in $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we just set $g = \sigma_t^{-1} \circ f_1 \circ \sigma_t$;
- in $\sigma_t^{-1}\big(f_1^{-1}(\sqcup_i B_\rho(v_i))\big) \setminus A^t$ we use the paths γ_α. More precisely, let N be a connected component of $f_1^{-1}(B_\rho(v_i))$ and let $x \in N$ such that $f_1(x) = v_i$. For any edge $\alpha \in N$ emanating from x we define $g(\sigma_t^{-1}(\alpha))$ by mapping linearly[11] $\sigma_t^{-1}(\alpha)$ to the path given by the concatenation of $\beta_Y = \sigma_t^{-1}(f_1(\alpha))$ and γ_α. Note that $g|_{\sigma_t^{-1}(\alpha)} = \mathrm{PL}(g|_{\sigma_t^{-1}(\alpha)})$.
- in A^t we set $g = \varphi$;

finally, we set

$$g_t = \mathrm{opt}(\mathrm{PL}(g))$$

where PL-ization and optimization are made with respect to the metric structure of Y_t. We now estimate the Lipschitz constant of g. Clearly

$$\lambda(Y_t) = \mathrm{Lip}(g_t) \leq \mathrm{Lip}(g).$$

In $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we have $g = \sigma_t^{-1} \circ f_1 \circ \sigma_t$. Then

$$\mathrm{Lip}(g) \leq \mathrm{Lip}(\sigma_t^{-1})\,\mathrm{Lip}(f_1)\,\mathrm{Lip}(\sigma_t).$$

Since on edges of $Y \setminus A$ the map σ_t is just a rescaling of edge-lengths, for any $\varepsilon_2 > 0$ there is $t_{\varepsilon_2} > 0$ such that $\forall t < t_{\varepsilon_2}$

$$\mathrm{Lip}(\sigma_t) < 1 + \varepsilon_2 \qquad \mathrm{Lip}(\sigma_t^{-1}) < 1 + \varepsilon_2 \tag{5}$$

hence, by (3), and by setting $(1 + \varepsilon_2)^2(1 + \varepsilon_1) = 1 + \varepsilon_3$ we have

$$\mathrm{Lip}(g) \leq (1 + \varepsilon_2)^2 \lambda(X)(1 + \varepsilon_1) = (1 + \varepsilon_3)\lambda(X). \tag{6}$$

[11] I.e. at constant speed

Now, let N be a connected component of $f_1^{-1}(\sqcup_i B_\rho(v_i))$. Let $x \in N$ such that $f_1(x) = v_i$ and let α be an edge of N emanating from x. By definition g is linear on $\sigma_t^{-1}(\alpha)$, thus in order to estimate its Lipschitz constant we need to know only the lengths of $\sigma_t^{-1}(\alpha)$ and its image. We have $L_X(f_1(\alpha)) = \rho$ and therefore

$$\rho \leq \mathrm{Lip}(f_1) L_X(\alpha) \qquad L_X(\alpha) = L_X(\sigma_t(\sigma_t^{-1}(\alpha))) \leq \mathrm{Lip}(\sigma_t) L_{Y_t}(\sigma_t^{-1}(\alpha))$$

whence, by (5) and (6), we obtain

$$L_{Y_t}(\sigma_t^{-1}(\alpha)) \geq \frac{L_X(\alpha)}{\mathrm{Lip}(\sigma_t)} > \frac{\rho}{(1+\varepsilon_2)\,\mathrm{Lip}(f_1)} > \frac{\rho}{\lambda(X)(1+\varepsilon_1)(1+\varepsilon_2)}.$$

Since γ_α is the same loop in A for every t, its length in A^t depends linearly on t, namely here is a constant C_α such that

$$L_{Y_t}(\gamma_\alpha) = C_\alpha t$$

whence, setting $C = \max_\alpha C_\alpha$,

$$\begin{aligned}
\mathrm{Lip}(g|_{\sigma_t^{-1}(\alpha)}) &\leq \frac{L_{Y_t}(\sigma_t^{-1}(f_1(\alpha)) + L_{Y_t}(\gamma_\alpha)}{L_{Y_t}(\sigma_t^{-1}(\alpha))} \leq \frac{\mathrm{Lip}(\sigma_t^{-1})\rho + tC_\alpha}{L_{Y_t}(\sigma_t^{-1}(\alpha))} \\
&< ((1+\varepsilon_2)\rho + tC_\alpha)\frac{\lambda(X)(1+\varepsilon_1)(1+\varepsilon_2)}{\rho} \\
&= \lambda(x)\big[(1+\varepsilon_3) + \frac{(1+\varepsilon_1)(1+\varepsilon_2) + tC_\alpha}{\rho}\big] \\
&< \lambda(x)\big[(1+\varepsilon_3) + \frac{(1+\varepsilon_3) + tC}{\rho}\big]
\end{aligned}$$

Therefore $\forall \varepsilon_4 > 0 \exists t_{\varepsilon_4} > 0$ such that $\forall t < t_{\varepsilon_4}$, for any α as above we have

$$\mathrm{Lip}(g|_{\sigma_t^{-1}(\alpha)}) < \lambda(X)(1+\varepsilon_4). \tag{7}$$

Finally, on A^t we have $g = \varphi$ and so $\mathrm{Lip}(g|_{A^t}) = \mathrm{Lip}(\varphi)$. Since by (2) $\lambda(X) \leq \lambda(Y_t)(1+\varepsilon_0)$, by putting together (4), (6), and (7) we have that for any $\varepsilon_5 > 0$ there is $t_{\varepsilon_5} > 0$ such that for any $t < t_{\varepsilon_5}$ we have

$$\mathrm{Lip}(g) \leq \lambda(Y_t)(1+\varepsilon_5)$$

Since g_t is optimal $\mathrm{Lip}(g_t) = \lambda(Y_t)$ and by Theorem 3.15

$$d_\infty(g_t, g) < \mathrm{vol}(Y_t)(\mathrm{Lip}(g) - \lambda(Y_t)) < \mathrm{vol}(Y_t)\lambda(Y_t)\varepsilon_5.$$

We now estimate

$$d_\infty(\sigma_t \circ g, f_1 \circ \sigma_t).$$

In $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we have $g = \sigma_t^{-1} \circ f_1 \circ \sigma_t$ so here the distance is zero. On A^t, since $g(A) = A$, for any i there is j such that we have $\sigma_t(g(A_i)) = \sigma_t(A_j) = v_j = f(v_i)$, hence also in A^t the distance is zero. Finally, let N be a connected component of $f_1^{-1}(\sqcup_i B_\rho(v_i))$. Let $x \in N$ such that $f_1(x) = v_i$ and let α be an edge of N emanating from x.

The path $g(\sigma_t^{-1}(\alpha))$ is given by the concatenation of $\sigma_t^{-1}(f_1(\alpha))$ with γ_α. The latter is collapsed by σ_t, and the image of the former is just $f_1(\alpha) = f_1 \circ \sigma_t(\sigma_t^{-1}(\alpha))$. Since the length of γ_α in A^t is bounded by tC we have that

$$d_\infty(\sigma_t \circ g, f_1 \circ \sigma_t) \to 0 \qquad \text{as } t \to 0.$$

In particular $\forall \varepsilon_6 \exists t_{\varepsilon_6}$ such that $\forall t < t_{\varepsilon_6}$ we have

$$d_\infty(\sigma_t \circ g, f_1 \circ \sigma_t) < \varepsilon_6.$$

Finally,

$$\begin{aligned} & d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) \\ \leq\ & d_\infty(\sigma_t \circ g_t, \sigma_t \circ g) + d_\infty(\sigma_t \circ g, f_1 \circ \sigma_t) + d_\infty(f_1 \circ \sigma_t, f \circ \sigma_t) \\ \leq\ & \mathrm{Lip}(\sigma_t) d_\infty(g_t, \circ g) + \varepsilon_6 + d_\infty(f_1, f) \\ <\ & (1+\varepsilon_2)\,\mathrm{vol}(Y_t)\lambda(Y_t)\varepsilon_5 + \varepsilon_6 + \varepsilon_1 \end{aligned}$$

which is arbitrarily small for $t \to 0$. □

Lemma 12.4. *Let $X, Y \in \overline{\mathcal{O}(\Gamma)}$. Suppose that as graphs of groups, X is obtained from Y by collapsing a sub-forest $T = \sqcup T_i$ whose tree T_i each contains at most one non-free vertex. For $t \in [0,1]$ let $Y_t = (1-t)X + tY$ be a parametrization of the Euclidean segment from X to Y. Let $\sigma_t : Y_t \to X$ be the map obtained by collapsing T and by linearly rescaling the edges in $Y \setminus T$.*

Let $f : X \to X$ be an optimal map representing ϕ. Then for any $\varepsilon > 0$ there is $t_\varepsilon > 0$ such that $\forall 0 \leq t < t_\varepsilon$ there is an optimal map $g_t : Y_t \to Y_t$ representing ϕ such that

$$d_\infty(\sigma_t \circ g_t, f \circ \sigma_t) < \varepsilon.$$

Proof. The proof goes exactly as that of Lemma 12.1, and it is even simpler. As above T^t denote the scaled version of T. Let v_i be the vertex of X resulting from the collapse of T_i. The function λ is now continuous

$$\lambda(Y_t) \to \lambda(X)$$

as above, if f collapses some edge we fine $f_1 : X \to X$ a PL-map representing ϕ which collapses no edge and with

$$d_\infty(f, f_1) < \varepsilon_1 \quad \text{and} \quad \mathrm{Lip}(f_1) < \mathrm{Lip}(f)(1+\varepsilon_1) = \lambda(X)(1+\varepsilon_1).$$

We choose ρ so that $B_\rho(v_i)$ is star-shaped, the components of $f_1^{-1}(B_\rho(v_i))$ are star-shaped and contain a unique pre-image of v_i, and so that the components of $f_1^{-1}(B_\rho(v_i))$ and $f_1^{-1}(B_\rho(v_j))$ are pairwise disjoint. Finally we chose ρ small enough so that if $f(v_i) \notin \{v_j\}$, then $f(v_i) \notin \cup_j B_\rho(v_j)$.

For any i we choose a base vertex $x_i \in T_i$ which the non-free vertex of T_i if any. For any $x \in X$ such that $f_1(x) = v_i$ and for any edge α in $f_1^{-1}(B_\rho(v_i))$ incident to x, let γ_α be the unique embedded path connecting $\sigma_t^{-1}(f_1(\gamma_\alpha))$ to x_i. We define $g : Y_t \to Y_t$ as follows:

- in $\sigma_t^{-1}\big(X \setminus f_1^{-1}(\sqcup_i B_\rho(v_i))\big)$ we just set $g = \sigma_t^{-1} \circ f_1 \circ \sigma_t$;
- in $\sigma_t^{-1}\big(f_1^{-1}(\sqcup_i B_\rho(v_i))\big) \setminus T^t$ we use the paths γ_α. More precisely, let N be a connected component of $f_1^{-1}(B_\rho(v_i))$ and let $x \in N$ such that $f_1(x) = v_i$. For any edge $\alpha \in N$ emanating from x we define $g(\sigma_t^{-1}(\alpha))$ by mapping linearly[12] $\sigma_t^{-1}(\alpha)$ to the path given by the concatenation of $\sigma_t^{-1}(f_1(\alpha))$ and γ_α. Note that $g|_{\sigma_t^{-1}(\alpha)} = \mathrm{PL}(g|_{\sigma_t^{-1}(\alpha)})$.
- in the components T_i^t so that $f_1(v_i) = v_j$, we set $g(T_i^t) = x_j$;

finally we set $g_t = \mathrm{opt}(\mathrm{PL}(g))$. The estimates on Lipschitz constants and distances now follow exactly as in the proof of Lemma 12.1. □

References

[1] Yael Algom-Kfir, *Strongly contracting geodesics in outer space*, Geom. Topol. **15** (2011), no. 4, 2181–2233. MR 2862155

[2] Yael Algom-Kfir and Mladen Bestvina, *Asymmetry of outer space*, Geom. Dedicata **156** (2012), 81–92. MR 2863547

[3] Mladen Bestvina, *A Bers-like proof of the existence of train tracks for free group automorphisms*, Fund. Math. **214** (2011), no. 1, 1–12. MR 2845630

[4] Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51. MR 1147956

[5] Stefano Francaviglia and Armando Martino, *Stretching factors, metrics and train tracks for free products*, Illinois J. Math. **59** (2015), no. 4, 859-899.

[6] Stefano Francaviglia and Armando Martino, *Metric properties of outer space*, Publ. Mat. **55** (2011), no. 2, 433–473. MR 2839451

[7] ———, *The isometry group of outer space*, Adv. Math. **231** (2012), no. 3-4, 1940–1973. MR 2964629

[8] S. M. Gersten, *Addendum: "On fixed points of certain automorphisms of free groups"*, Proc. London Math. Soc. (3) **49** (1984), no. 2, 340–342. MR 748994

[9] ———, *On fixed points of certain automorphisms of free groups*, Proc. London Math. Soc. (3) **48** (1984), no. 1, 72–90. MR 721773

[10] ———, *Fixed points of automorphisms of free groups*, Adv. in Math. **64** (1987), no. 1, 51–85. MR 879856

[11] Vincent Guirardel and Gilbert Levitt, *The outer space of a free product*, Proc. Lond. Math. Soc. (3) **94** (2007), no. 3, 695–714. MR 2325317

[12] Ilya Kapovich, *Detecting fully irreducible automorphisms: a polynomial time algorithm*, https://arxiv.org/abs/1609.03820.

[13] Ilya Kapovich, *Algorithmic detectability of iwip automorphisms*, Bull. Lond. Math. Soc. **46** (2014), no. 2, 279–290. MR 3194747

[14] Jérôme E. Los, *On the conjugacy problem for automorphisms of free groups*, Topology **35** (1996), no. 3, 779–808, With an addendum by the author. MR 1396778

[15] Martin Lustig, *Conjugacy and centralizers for iwip automorphisms of free groups*, Geometric group theory, Trends Math., Birkhäuser, Basel, 2007, pp. 197–224. MR 2395795

[16] Sebastian Meinert, *The Lipschitz metric on deformation spaces of G-trees*, Algebr. Geom. Topol. **15** (2015), no. 2, 987–1029. MR 3342683

[17] John R. Stallings, *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565. MR 695906

[12] I.e. at constant speed

[18] E. Ventura, *Fixed subgroups in free groups: a survey*, Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001), Contemp. Math., vol. 296, Amer. Math. Soc., Providence, RI, 2002, pp. 231–255. MR 1922276